

Galois theory, graphs and free groups

Brent Everitt*

Abstract. A self-contained exposition is given of the topological and Galois-theoretic properties of the category of combinatorial 1-complexes, or graphs, very much in the spirit of Stallings [20]. A number of classical, as well as some new results about free groups are derived.

Introduction

This paper is about the interplay between graphs, free groups, and their subgroups, a subject with a long history that can be broadly divided into two schools. The first is combinatorial, where graphs, and particularly finite graphs, provide a intuitively convenient way of picturing some aspects of the theory of free groups, as in for example [10, 11, 19, 21, 22].

The other approach is to treat graphs and their mappings as topological objects, a point of view with its origins from the very beginnings of combinatorial group theory, and resurrected in [20] (see also [2, 6, 13]). This is the philosophy we take, but we differ from these earlier papers in that we place centerstage the theory of coverings of arbitrary graphs, rather than coverings being merely a prelude to immersions of finite graphs. The first section sets up the topological preliminaries, §2 formulates the well known connection between subgroups of free groups and coverings of graphs in a Galois-theoretic setting, while §§3-4 focus on the graph-theoretic implications of finitely generated subgroups.

1. The topology of graphs

This section is all very “Stallings-esque” [20], with much of the material in §§1.1-1.8 well known. General references are [2, 3, 6, 17, 18, 20]. In §1.9 we deal with coverings, with a mixture of well known and some (minor) new results; §1.10 introduces the lattice of intermediate coverings of a cover.

1.1. Graphs

A *combinatorial 1-complex* or *graph* [6, §1.1] is an arbitrary set Γ with an involutory map $^{-1} : \Gamma \rightarrow \Gamma$ and an idempotent map $s : \Gamma \rightarrow V_\Gamma$, (ie: $s^2 = s$) where V_Γ is the set of fixed points of $^{-1}$. Thus a graph has *vertices* V_Γ , and *edges* $E_\Gamma := \Gamma \setminus V_\Gamma$ with (i). $s(v) = v$ for all $v \in V_\Gamma$; (ii). $v^{-1} = v$ for all $v \in V_\Gamma$, $e^{-1} \in E_\Gamma$ and $e^{-1} \neq e = (e^{-1})^{-1}$ for all $e \in E_\Gamma$. Indeed, these two can be taken as a more transparent, but less elegant, definition. We will use both interchangeably.

The edge e has *start vertex* $s(e)$ and *terminal vertex* $t(e) := s(e^{-1})$; an *arc* is an edge/inverse edge pair, and an *orientation* for Γ is a set \mathcal{O} of edges containing exactly one edge from each arc. Write \bar{e} for the arc containing the edge e (so that $e^{-1} = \bar{e}$). A pointed graph is a pair $\Gamma_v := (\Gamma, v)$ for $v \in \Gamma$ a vertex.

BRENT EVERITT: Department of Mathematics, University of York, York YO10 5DD, United Kingdom. e-mail: bje1@york.ac.uk

* Some of the results of this paper were obtained while the author was visiting the Department of Mathematics, University of Adelaide, Australia.

The graph Γ is *finite* when V_Γ is finite and *locally finite* when the set $s^{-1}(v)$ is finite for every $v \in V_\Gamma$. The cardinality of the set $s^{-1}(v)$ is the *valency* ∂v of the vertex v . A *path* is a finite sequence of edges, mutually incident in the obvious sense; similarly we have *closed* paths and *trivial* paths (consisting of a single vertex). Γ is *connected* if any two vertices can be joined by a path. The connected component of Γ containing the vertex v consists of those vertices for which there is a path connecting them to v , together with all their incident edges.

A *map* of graphs is a set map $g : \Gamma \rightarrow \Lambda$ with $g(V_\Gamma) \subseteq V_\Lambda$, such that the diagram commutes, where σ_Γ is one of the s or $^{-1}$ maps for Γ , and σ_Λ similarly, ie: $gs_\Gamma(x) = s_\Lambda g(x)$ and $g(x^{-1}) = g(x)^{-1}$. These are combinatorial versions of continuity: if Γ is connected then $g(\Gamma) \subset \Lambda$ is connected. A map is *dimension preserving* if $g(E_\Gamma) \subseteq E_\Lambda$. These maps of graphs allow one to squash edges down to vertices as in [6], rather than the more rigid maps of [18, 20]. The pay off is that the quotient construction below is more useful. A map $g : \Gamma_v \rightarrow \Lambda_u$ of pointed graphs is a map $g : \Gamma \rightarrow \Lambda$ with $g(v) = u$.

A map $g : \Gamma \rightarrow \Lambda$ is a *homeomorphism* if it is dimension preserving and is a bijection on the vertex and edge sets, in which case the inverse set map is a dimension preserving map of graphs $g^{-1} : \Lambda \rightarrow \Gamma$, and hence a homeomorphism.

The set of self homeomorphisms $\Gamma \rightarrow \Gamma$ forms a group $\text{Homeo}(\Gamma)$, and a group action $G \xrightarrow{\varphi} \text{Homeo}(\Gamma)$ is said to *preserve orientation* iff there is an orientation \mathcal{O} for Γ with $\varphi(g)(\mathcal{O}) = \mathcal{O}$ for all $g \in G$. The action of G is said to be *without inversions* iff $\varphi(g)(e) \neq e^{-1}$ for all edges e and for all $g \in G$. It is easy to see that G preserves orientation if and only if it acts without inversions. G acts *freely* iff the action is free on the vertices, ie: for any $g \in G$ and v a vertex, $\varphi(g)(v) = v$ implies that g is the identity element. If G acts freely and orientation preservingly, then the action is free on the edges too.

A *subgraph* is a subset $\Lambda \subset \Gamma$, such that the maps s and $^{-1}$ give a graph when restricted to Λ . Equivalently, it is a graph mapping $\Lambda \hookrightarrow \Gamma$ that is a homeomorphism onto its image. The *coboundary* $\delta\Lambda$ of a subgraph consists of those edges $e \in \Gamma$ with $s(e) \in \Lambda$ and $t(e) \notin \Lambda$ (equivalently, it is those edges $e \in \Gamma$ with $sq(e)$ the vertex $q(\Lambda)$ in the quotient complex Γ/Λ , where $q : \Gamma \rightarrow \Gamma/\Lambda$ is the quotient mapping as below).

An *elementary homotopy* of a path, $e_1 \dots e_i e_{i+1} \dots e_k \leftrightarrow e_1 \dots e_i (ee^{-1}) e_{i+1} \dots e_k$ inserts or deletes a *spur*: a path that consecutively traverses both edges of an arc \bar{e} . Two paths are (freely) *homotopic* iff there is a finite sequence of elementary homotopies taking one to the other. Paths homotopic to a trivial path are said to be *homotopically trivial*. It is easy to see that two homotopic paths have the same start and terminal vertices (and thus homotopically trivial paths are necessarily closed) and that homotopy is an equivalence relation on the paths with common endpoints.

The *trivial graph* has a single vertex and no edges. The *real line* graph \mathcal{R} has vertices $V_{\mathcal{R}} = \{v_k\}_{k \in \mathbb{Z}}$, edges $E_{\mathcal{R}} = \{e_k^{\pm 1}\}_{k \in \mathbb{Z}}$ and $s(e_k) = v_k, s(e_k^{-1}) = v_{k+1}$.

1.2. Quotients

A *quotient relation* is an equivalence relation \sim on Γ such that

$$(i). x \sim y \Rightarrow s(x) \sim s(y) \text{ and } x^{-1} \sim y^{-1}, \quad (ii). x \sim x^{-1} \Rightarrow [x] \cap V_\Gamma \neq \emptyset,$$

where $[x]$ is the equivalence class under \sim of x .

Proposition 1. *If \sim is a quotient relation on a graph Γ then the quotient graph Γ/\sim has vertices the equivalence classes $[v]$ for $v \in V_\Gamma$, edges the classes $[e]$ for $e \in E_\Gamma$ with $[e] \cap V_\Gamma = \emptyset$, $[x]^{-1} = [x^{-1}]$, and $s_{\Gamma/\sim}[x] = [s_\Gamma(x)]$. Moreover, the quotient map $q : \Gamma \rightarrow \Gamma/\sim$ given by $q(x) = [x]$ is a map of graphs (and so in particular, if Γ is connected then Γ/\sim is connected).*

Let $\Lambda_i \hookrightarrow \Gamma, (i \in I)$, be a set of mutually disjoint subgraphs and define \sim on Γ by $x \sim y$ iff $x = y$ or both x and y lie in the same Λ_i . Write Γ/Λ_i for Γ/\sim , the *quotient of Γ by the family of subgraphs Λ_i* . It is what results by squashing each Λ_i to a distinct vertex. In particular,

if the family consists of a single subgraph $\Lambda \hookrightarrow \Gamma$, we have the quotient Γ/Λ . The reader should be wary of the difference between the quotients Γ/Λ_i and Γ/Λ , for $\Lambda = \coprod \Lambda_i$ the union of the disjoint subgraphs.

If \sim is the equivalence relation on Γ consisting of the orbits of an action by a group G , then \sim is a quotient relation on Γ if and only if the group action is orientation preserving. In this case we may form the quotient complex $\Gamma/G := \Gamma/\sim$.

1.3. Trees

A path in a graph is *reduced* when it contains no spurs; by removing spurs, any two vertices in the same component can be joined by a reduced path.

It is easily proved that for any vertices u and v of a graph Γ , there are ≤ 1 reduced paths between them if and only if any closed non-trivial path in Γ contains a spur (equivalently, any closed path is homotopic to the trivial path based at one of its vertices). A graph satisfying any of these equivalent conditions is called a *forest*; a connected forest is a *tree*.

If Γ is a finite graph with $\partial v \geq 2$ for every vertex v , then it can be shown that Γ contains a homotopically non-trivial closed path. Hence if T is a finite tree, then $|E_T| = 2(|V_T| - 1)$.

A *spanning forest* is a subgraph $\Phi \hookrightarrow \Gamma$ that is a forest and contains all the vertices of Γ (ie: $V_\Phi = V_\Gamma$). It is well known that spanning trees can always be constructed for connected Γ .

Proposition 2. *Let $T_i \hookrightarrow \Gamma$ be a family of mutually disjoint trees in a connected graph Γ . Then there is a spanning tree $T \hookrightarrow \Gamma$ containing the T_i as subgraphs, and such that if $q : \Gamma \rightarrow \Gamma/T_i$ is the quotient map, then $q(T)$ is a spanning tree for Γ/T_i .*

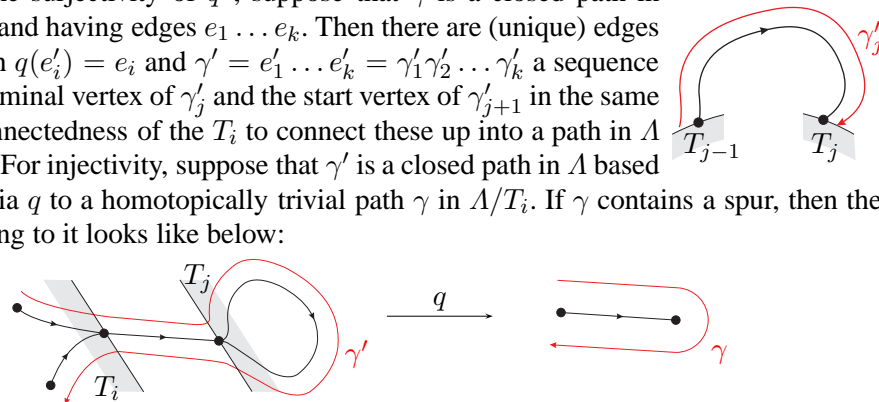
In particular, any spanning forest for Γ can be extended to a spanning tree. For the proof, take T to be $q^{-1}(T')$ for some spanning tree T' of the (connected) graph Γ/T_i .

1.4. The fundamental group

The *fundamental group* $\pi_1(\Gamma, v)$ is the usual group of homotopy classes $[\gamma]$ of closed paths γ at the vertex $v \in \Gamma$ (ie: equivalence classes under the homotopy relation) with product $[\gamma_1][\gamma_2] = [\gamma_1\gamma_2]$. If Φ is a forest, then $\pi_1(\Phi, v)$ is trivial for any vertex v and conversely, if Γ connected has $\pi_1(\Gamma, v)$ trivial for some (hence every) vertex v , then Γ is a tree. A connected graph with trivial fundamental group is *simply connected*. A map $g : \Gamma_v \rightarrow \Lambda_u$ of graphs induces a group homomorphism $g^* : \pi_1(\Gamma, v) \rightarrow \pi_1(\Lambda, u)$ by $g^*[\gamma] = [g(\gamma)]$ and this satisfies the usual functoriality properties: $(\text{id})^* = \text{id}$ and $(gf)^* = g^*f^*$.

Proposition 3. *If $T_i \hookrightarrow \Lambda$ is a family of mutually disjoint trees, $v \in T \in \{T_i\}$ a vertex, and $q : \Lambda \rightarrow \Lambda/T_i$ the quotient map, then $q^* : \pi_1(\Lambda, v) \rightarrow \pi_1(\Lambda/T_i, q(v))$ is an isomorphism.*

Proof. The key to the proof is that the quotient map q is essentially just the identity map outside of the T_i . To see the surjectivity of q^* , suppose that γ is a closed path in Λ/T_i based at $q(v)$ and having edges $e_1 \dots e_k$. Then there are (unique) edges e'_1, \dots, e'_k in Λ with $q(e'_i) = e_i$ and $\gamma' = e'_1 \dots e'_k = \gamma'_1 \gamma'_2 \dots \gamma'_k$ a sequence of paths with the terminal vertex of γ'_j and the start vertex of γ'_{j+1} in the same tree T_j . Use the connectedness of the T_i to connect these up into a path in Λ mapping via q to γ . For injectivity, suppose that γ' is a closed path in Λ based at v and mapping via q to a homotopically trivial path γ in Λ/T_i . If γ contains a spur, then the section of γ' mapping to it looks like below:



Thus a sequence of elementary homotopies reducing γ to the trivial path in Λ/T_i can be mirrored by homotopies in Λ that reduce γ' to a closed path completely contained in T . As T is simply connected, this path can in turn be homotoped to the trivial path. Thus, only homotopically trivial paths can be sent by q to homotopically trivial paths, so q^* is injective. \square

Fix a spanning tree $T \hookrightarrow \Gamma$, choose an edge e from each arc of Γ , and consider the homotopy class of the path through T from v to $s(e)$, traverses e and travels back through T to v . Then *Schreier generators* for $\pi_1(\Gamma, v)$ are the homotopy classes of such paths arising from the arcs omitted by T .

1.5. Homology

Fix an orientation \mathcal{O} for Γ , and always write arcs in the form \bar{e} for $e \in \mathcal{O}$, and paths in the form $\gamma = e_1^{\varepsilon_1} \dots e_k^{\varepsilon_k}$ with $e_i \in \mathcal{O}$ and $\varepsilon_i = \pm 1$. Let $\mathbb{Z}[V_\Gamma]$ and $\mathbb{Z}[\text{arcs}]$ be the free abelian groups on the vertices and arcs of Γ (alternatively, one can take $\mathbb{Z}[E_\Gamma]$ and then pass to the quotient $\mathbb{Z}[E_\Gamma]/\langle e + e^{-1} = 0 \rangle$; we prefer the more concrete version). Define the boundary of an arc \bar{e} to be $\partial(\bar{e}) = t(e) - s(e) \in \mathbb{Z}[V_\Gamma]$, and for $\sum n_i \bar{e}_i \in \mathbb{Z}[\text{arcs}]$, let $\partial(\sum n_i \bar{e}_i) = \sum n_i \partial(\bar{e}_i)$. Then ∂ is a group homomorphism $\partial : \mathbb{Z}[\text{arcs}] \rightarrow \mathbb{Z}[V_\Gamma]$, and the *homology* of Γ is the pair of abelian groups

$$H_1(\Gamma) = \ker \partial \text{ and } H_0(\Gamma) = \text{coker } \partial,$$

(ie: $H_0(\Gamma) = \mathbb{Z}[V_\Gamma]/\text{im } \partial$).

By following the proofs in the topological category, one can show the standard homological facts: $H_0(\Gamma)$ is free abelian on the connected components of Γ ; if Γ is single vertexed, then H_1 is free abelian on the arcs. If $\gamma = e_1^{\varepsilon_1} e_2^{\varepsilon_2} \dots e_k^{\varepsilon_k}$ is a closed path at v then $\partial(\sum \varepsilon_i \bar{e}_i) = 0$, and the Hurewicz map sending γ to $\sum \varepsilon_i \bar{e}_i$ is well defined upto homotopy, thus, for Γ connected, a surjective homomorphism $\pi_1(\Gamma, v) \rightarrow H_1(\Gamma)$ with kernel the commutator subgroup of $\pi_1(\Gamma, v)$.

In particular, $H_1(\Gamma)$ is the abelianisation of $\pi_1(\Gamma, v)$, so that if Γ_1, Γ_2 are connected graphs with $\pi_1(\Gamma_1, v_1) \cong \pi_1(\Gamma_2, v_2)$ then $H_1(\Gamma_1) \cong H_1(\Gamma_2)$.

1.6. Rank and spines

Graph homology provides an important invariant for graphs:

Proposition 4 (rank an invariant). *Let $T \hookrightarrow \Gamma$ be a spanning tree for Γ connected. Then $H_1(\Gamma)$ is free abelian with basis the set of arcs of Γ omitted by T .*

Thus the cardinality of the set of omitted arcs is equal to $\text{rk}_{\mathbb{Z}} H_1(\Gamma)$ and independent of T (this can also be shown directly without recourse to homology). Define the *rank* $\text{rk } \Gamma$ of Γ connected to be $\text{rk}_{\mathbb{Z}} H_1(\Gamma)$, or the cardinality of the set of arcs omitted by a spanning tree.

Proof (of Proposition 4). We have $\pi(\Gamma, v) \cong \pi_1(\Gamma/T, q(v))$ by Proposition 3, hence $H_1(\Gamma) \cong H_1(\Gamma/T)$, with Γ/T single vertexed, hence this final homology free abelian on its arcs, ie: free abelian on the arcs of Γ omitted by T . \square

If Γ is finite, locally finite, connected, then, $2(\text{rk } \Gamma - 1) = |E_\Gamma| - 2|V_\Gamma|$ by §1.3; clearly, $\text{rk } \Gamma = 0$ if and only if Γ is a tree. If Γ a connected graph and $T_i \hookrightarrow \Gamma$ a set of mutually disjoint trees, then $\text{rk } \Gamma = \text{rk } \Gamma/T_i$ (this follows either from Proposition 3 using $\text{rk} = \text{rk } H_1$ or by Proposition 2 using rank the number of arcs omitted by a spanning tree).

If Λ is a connected graph and v a vertex, then the *spine* $\hat{\Lambda}_v$ of Λ at v , is defined to be the union in Λ of all closed reduced paths starting at v . Stallings and others use core graphs; we have followed [13].

Lemma 1. (i). $\hat{\Lambda}_v$ is connected with $\text{rk } \hat{\Lambda}_v = \text{rk } \Lambda$. (ii). If $u \in \hat{\Lambda}_v$, then every closed reduced path starting at u is contained in $\hat{\Lambda}_v$. (iii). Spines are topological invariants, ie: a homeomorphism $f : \Lambda_u \rightarrow \Delta_v$ restricts to a homeomorphism $\hat{\Lambda}_u \rightarrow \hat{\Delta}_v$.

Proof. The connectedness is immediate. If T is a spanning tree for Λ and e an edge not in T , then e is contained in the spine $\widehat{\Lambda}_v$, for, the closed path obtained by traversing the reduced path through T from v to $s(e)$, across e and back via the reduced path through T is reduced. The rank assertion follows. For part (ii), let μ be a closed reduced path at u and γ a reduced path in the spine from v to u . Then $\gamma = \gamma'_1 \gamma_1 = \gamma'_2 \gamma_2$ and $\mu = \gamma_1^{-1} \mu' \gamma_2$ with $\gamma'_1 \mu' (\gamma'_2)^{-1}$ reduced at v , hence in the spine. A homeomorphism sends closed reduced paths to closed reduced paths (compare with Proposition 7 and the maps of §1.9), hence $f(\widehat{\Lambda}_v) \subset \widehat{\Delta}_u$, and the converse similarly using f^{-1} . \square

1.7. Pushouts

These are important examples of quotients. Let Λ_1, Λ_2 and Δ be graphs and $g_i : \Delta \rightarrow \Lambda_i$ maps of graphs. Let \sim on the disjoint union $\Lambda_1 \amalg \Lambda_2$ be the equivalence relation generated by the $x \sim y$ iff there is a $z \in \Delta$ with $x = g_1(z)$ and $y = g_2(z)$. Thus, $x \sim y$ iff there are x_0, x_1, \dots, x_k with $x_0 = x$ and $x_k = y$, and for each j , there is $z \in \Delta$ with $x_j = g_1(z)$, $x_{j+1} = g_2(z)$. If \sim is a quotient relation then call the quotient $\Lambda_1 \amalg \Lambda_2 / \sim$ the *pushout* of the data $g_i : \Delta \rightarrow \Lambda_i$, denoted $\Lambda_1 \amalg_{\Delta} \Lambda_2$.

Given graphs and maps as above, the pushout cannot always be formed, precisely because the quotient cannot always be formed. Stallings [20, page 552] shows that if the g_i are dimension preserving, then the pushout exists if and only if there are orientations $\mathcal{O}, \mathcal{O}_i$ for Δ, Λ_i with $g_i(\mathcal{O}) \subseteq \mathcal{O}_i$. Thus in particular, if the graphs $g_1(\Delta)$ and $g_2(\Delta)$ are disjoint, then the pushout can always be formed.

Define $t_i : \Lambda_i \rightarrow \Lambda_1 \amalg_{\Delta} \Lambda_2$ to be the compositions $\Lambda_i \hookrightarrow \Lambda_1 \amalg \Lambda_2 \rightarrow \Lambda_1 \amalg \Lambda_2 / \sim$ of the inclusion of Λ_i in the disjoint union and the quotient map.

Proposition 5. *If $\Delta \neq \emptyset$ and the Λ_i are connected then the pushout is connected, and the maps t_i make the diagram on the left commute.*

$$\begin{array}{ccc}
 \Delta & \xrightarrow{g_2} & \Lambda_2 \\
 \downarrow g_1 & & \downarrow t_2 \\
 \Lambda_1 & \xrightarrow{t_1} & \Lambda_1 \amalg_{\Delta} \Lambda_2
 \end{array}
 \qquad
 \begin{array}{ccc}
 \Delta & \xrightarrow{g_2} & \Lambda_2 \\
 \downarrow g_1 & & \downarrow t_2 \\
 \Lambda_1 & \xrightarrow{t_1} & \Lambda_1 \amalg_{\Delta} \Lambda_2 \\
 & \searrow t'_1 & \downarrow t'_2 \\
 & & B
 \end{array}$$

Moreover the pushout is universal in that if B, t'_1, t'_2 are a graph and maps making such a square commute, then there is a map $\Lambda_1 \amalg_{\Delta} \Lambda_2 \rightarrow B$ making the diagram above right commute.

These properties can of course be taken as an alternative, categorical definition of the pushout, with uniqueness following from the universality and the usual formal nonsense. If the $g_i : \Delta_v \rightarrow \Lambda_{u_i}$ are pointed maps, and $q : \Lambda_1 \amalg \Lambda_2 \rightarrow \Lambda_1 \amalg \Lambda_2 / \sim$ the quotient map, then $q(u_1) = q(u_2) = u$ (say), and we have a pointed version of Proposition 5, involving the pointed pushout $(\Lambda_1 \amalg_{\Delta} \Lambda_2)_u$.

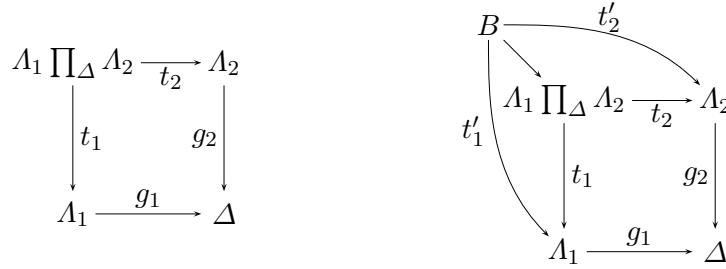
Many of the quotient constructions from topology (eg: cone, suspension, ...) can be expressed as some kind of pushout or other, but we content ourselves with the following: let Δ be a graph with $E_{\Delta} = \emptyset$ and the $g_i : \Delta \hookrightarrow \Lambda_i$ homeomorphisms onto their images (ie: injections on the vertices of Δ). The resulting pushout (which always exists), the *wedge sum* $\Lambda_1 \vee_{\Delta} \Lambda_2$, is the result of identifying the vertices of copies of Δ in the Λ_i . If the Λ_i coincide ($= \Lambda$ say) with maps $\Delta \rightrightarrows \Lambda$, then we write $\bigvee_{\Delta} \Lambda$.

If Δ is the trivial graph, T a tree and $g_i : \Delta \rightrightarrows T$ distinct maps, then the wedge sum $\bigvee_{\Delta} T$ has a non-trivial reduced closed path that is unique upto cyclic reordering. Thus, by removing a single arc from $\bigvee_{\Delta} T$ we obtain a new tree.

1.8. Pullbacks

The categorical nature of the pushout construction (ie: Proposition 5) suggests a “co-” version: let Λ_1, Λ_2 and Δ be graphs and $g_i : \Lambda_i \rightarrow \Delta$ maps of graphs. The *pullback* $\Lambda_1 \prod_{\Delta} \Lambda_2$ has vertices (resp. edges) the $x_1 \times x_2$, $x_i \in V_{\Lambda_i}$ (resp. $x_i \in E_{\Lambda_i}$) such that $g_1(x_1) = g_2(x_2)$, and $s(x_1 \times x_2) = s(x_1) \times s(x_2)$, $(x_1 \times x_2)^{-1} = x_1^{-1} \times x_2^{-1}$. Taking Δ to be the trivial graph has the effect of removing the $g_1(x) = g_2(y)$ conditions and the result is the *product* $\Lambda_1 \prod \Lambda_2$. Thus the pullback $\Lambda_1 \prod_{\Delta} \Lambda_2$ is a subgraph of the product $\Lambda_1 \prod \Lambda_2$, but the product will have many more vertices and edges. Define maps $t_i : \Lambda_1 \prod_{\Delta} \Lambda_2 \rightarrow \Lambda_i$ to be the compositions $\Lambda_1 \prod_{\Delta} \Lambda_2 \hookrightarrow \Lambda_1 \prod \Lambda_2 \rightarrow \Lambda_i$, with the second map the projection $x_1 \times x_2 \mapsto x_i$.

Proposition 6. *The t_i are dimension preserving maps making the diagram below left commute,*



Moreover, the pullback is universal in that if B, t'_1, t'_2 are a graph and maps making such a square commute, then there is a map $B \rightarrow \Lambda_1 \prod_{\Delta} \Lambda_2$ making the diagram above right commute.

In general the pullback need not be connected. If the $g_i : \Lambda_{u_i} \rightarrow \Delta_v$ are pointed maps then $u_1 \times u_2$ is a vertex of the pullback, and we may consider the connected component containing $u_1 \times u_2$. Call this the *pointed pullback* $(\Lambda_1 \prod_{\Delta} \Lambda_2)_{u_1 \times u_2}$, and we then have a pointed version of Proposition 6. In most of our usages of the pullback construction, the graph Δ will be single vertexed, and so the vertex set will just be $V_{\Lambda_1} \times V_{\Lambda_2}$.

1.9. Coverings

A map $p : \Lambda \rightarrow \Delta$ of graphs is a *covering* iff (i). p preserves dimension; and (ii). for every vertex $v \in \Delta$, p is a bijection from the set of edges in Λ with start vertex v to the set of edges in Δ with start vertex $p(v)$. If $p(x) = y$, then one says that x *covers* y , and y *lifts* to x . The set of all lifts of the cell y , or the set of all cells covering y , is its *fiber* $\text{fib}_{\Lambda \rightarrow \Delta}(y)$.

Proposition 7 ([20, §4.1]). *Let $p : \Lambda_u \rightarrow \Delta_v$ be a covering.*

- (i). *If γ is a path in Δ starting at v then there is a path γ' in Λ starting at u and covering γ . Moreover, if γ_1, γ_2 are paths in Δ starting at v and covering the same path, then $\gamma_1 = \gamma_2$.*
- (ii). *A path in Λ covering a spur is itself a spur. Consequently, two paths in Λ covering homotopic paths are homotopic.*
- (iii). *If $g : \Gamma_w \rightarrow \Delta_v$ is a map then there is a map $f : \Gamma_w \rightarrow \Lambda_u$ with $g = fp$ if and only if $g^* \pi_1(\Gamma, w) \subset p^* \pi_1(\Lambda, u)$.*
- (iv). *$p^* : \pi_1(\Lambda, u) \rightarrow \pi_1(\Delta, v)$ is injective, and if u' is the terminal vertex of a path μ starting at u , then $p^* \pi_1(\Lambda, u) = gp^* \pi_1(\Lambda, u')g^{-1}$, where g is the homotopy class of $p(\mu)$.*

The path γ' in (i) is a lift of γ to u , such lifts being unique by (ii). The combination of these two is called *path lifting*, while (ii) is *spur-lifting* and *homotopy lifting*. In particular, the image under a covering of a reduced path is reduced (whereas, as spurs always map to spurs, the pre-image of a reduced path is reduced under any mapping). Part (iii) is a general lifting criterion that implies in particular that if γ a closed path at v then its homotopy class lies in $p^* \pi_1(\Lambda, u)$ if and only if there is a closed path μ at u with $p(\mu) = \gamma$. Part (iv) follows immediately from this and homotopy lifting.

Lemma 2. Let $p : \Lambda \rightarrow \Delta$ be a covering.

- (i). If Δ is connected then p maps the cells of Λ surjectively onto the cells of Δ .
- (ii). If Λ is connected then the fibers of any two cells of Δ have the same cardinality, called the degree, $\deg(\Lambda \rightarrow \Delta)$, of the covering.
- (iii). If Λ, Δ are connected and $\deg(\Lambda \rightarrow \Delta) = 1$, then the covering $\Lambda \rightarrow \Delta$ is a homeomorphism.

Proof. In (i), surjectivity on the vertices follows by path lifting and on the edges by definition. Path lifting gives a bijection in (ii) between the fibers of two vertices, and between the fiber of an edge and its start vertex. Part (iii) follows immediately from (i) and (ii). \square

From now on, *all* coverings will be maps between connected complexes unless stated otherwise.

Lemma 3. (i). Let $\Lambda \xrightarrow{q} \Gamma \xrightarrow{r} \Delta$ be maps with $p = rq$. If any two of p, q and r are coverings, then so is the third.

(ii). If a group G acts orientation preservingly and freely on Λ then the quotient map $q : \Lambda \rightarrow \Lambda/G$ is a covering.

Call the coverings $\Lambda \rightarrow \Gamma \rightarrow \Delta$ in (i) *intermediate* to the covering $\Lambda \rightarrow \Delta$. It follows from the comments following Proposition 7 that if $p : \Lambda_u \rightarrow \Delta_v$ and $r : \Gamma_x \rightarrow \Delta_v$, then $p^*\pi_1(\Lambda, u) \subset r^*\pi_1(\Gamma, x)$.

Proof (of Lemma 3). The freeness of the action in (ii) ensures the injectivity of q on the edges starting at a vertex of Λ . Part (i) is an easy exercise. \square

Proposition 8. Let Λ be a graph and $\Upsilon_1, \Upsilon_2 \hookrightarrow \Lambda$ subgraphs of the form,

$$\Lambda = \boxed{\Upsilon_1} \bullet \xrightarrow{e} \bullet \boxed{\Upsilon_2} \quad (\dagger)$$

- (i). If $p : \Lambda \rightarrow \Delta$ is a covering with Δ single vertexed, then the real line is a subgraph $\alpha : \mathcal{R} \hookrightarrow \Lambda$, with $\alpha(e_0) = e$ and $p\alpha(e_k) = p(e)$ for all $k \in \mathbb{Z}$.
- (ii). If Υ_1 is a tree, $p : \Lambda \rightarrow \Delta$, $r : \Gamma \rightarrow \Delta$ coverings, and $\alpha : \Upsilon_2 \hookrightarrow \Gamma$ a homeomorphism onto its image, then there is an intermediate covering $\Lambda \xrightarrow{q} \Gamma \xrightarrow{r} \Delta$.
- (iii). If $\Psi \rightarrow \Lambda$ is a covering and Υ_1 a tree, then Ψ also has the form (\dagger) for some subgraphs $\Upsilon'_1, \Upsilon'_2 \hookrightarrow \Psi$, with Υ'_1 a tree.

Proof. (i). Lift the edge $p(e)$ to the vertex $t(e)$ to get an edge e_1 of Λ . The form of Λ prohibits $t(e_1)$ from being any vertex of Υ_1 , except possibly $s(e)$, in which case $e_1 = e^{-1}$. But then $p(e)^{-1} = p(e^{-1}) = p(e_1) = p(e)$ a contradiction. Thus e_1 is an edge and $t(e_1)$ a vertex of Υ_2 , and if $t(e_1) = t(e)$, then the injectivity of p fails at this common vertex (as then both e^{-1} and e_1^{-1} start at $t(e_1)$ and cover $p(e)^{-1}$). We therefore have $t(e_1) \neq t(e)$ and this process can be continued inductively, giving the “positive” half of \mathcal{R} a subgraph of Υ_2 . The symmetry of Λ gives the negative half a subgraph of Υ_1 .

For (ii), it suffices, by part Lemma 3(i), to find a map $q : \Lambda \rightarrow \Gamma$ with $p = rq$. Let q coincide with α on Υ_2 . For any vertex of Υ_1 , take the reduced path to it from $t(e)$, project via p to Δ , and lift to $\alpha t(e) \in \Gamma$. The edges of Υ_1 (and e) are similar.

(iii). Let $v = t(e)$ and u be in the fiber of v via the covering $\Psi \rightarrow \Lambda$. Take a reduced path in $\Upsilon_1 \cup \{e\}$ from v to each vertex of this tree and lift to a path at u . Let Υ'_1 be the union in Ψ of these lifted paths. A closed path in Υ'_1 at u covers a closed path at v in Υ_1 , a tree, hence by spur-lifting, Υ'_1 is a tree. If e_1, e_2 are edges in the coboundary $\delta\Upsilon'_1$ then they cover edges in the coboundary $\delta\Upsilon_1$, ie: they cover e . A reduced path in Υ'_1 from $s(e_1)$ to $s(e_2)$ covers a reduced closed path in Υ_1 at $s(e)$. As this covered path must be trivial we get $s(e_1) = s(e_2)$, hence $e_1 = e_2$. Thus Υ'_1 has a single coboundary edge as required. \square

Proposition 9. Let $p : \Lambda \rightarrow \Delta$ be a covering and $T \hookrightarrow \Delta$ a tree. Then (i). $p^{-1}(T)$ a forest. (ii). If $T_i \hookrightarrow \Lambda$, $(i \in I)$ are the connected components of $p^{-1}(T)$, then p maps each T_i homeomorphically onto T . (iii). There is an induced covering $\Lambda/T_i \rightarrow \Delta/T$ making the diagram,

$$\begin{array}{ccc} \Lambda & \longrightarrow & \Lambda/T_i \\ \downarrow & & \downarrow \\ \Delta & \longrightarrow & \Delta/T \end{array}$$

commute (where the horizontal maps are the quotients) and such that $\deg(\Lambda/T_i \rightarrow \Delta/T) = \deg(\Lambda \rightarrow \Delta)$.

This procedure is independent of the tree T : if $T' \hookrightarrow \Delta$ another tree such that there is a homeomorphism $\alpha : \Delta/T \rightarrow \Delta/T'$ with $\alpha q = q'$ for $q, q' : \Delta \rightarrow \Delta/T, \Delta/T'$ the quotient maps, then by Proposition 7(iv), there is a homeomorphism $\Lambda/T_i \rightarrow \Lambda/T'_i$. Typically we will take T, T' to be spanning trees, so that $\Delta/T, \Delta/T'$ are single vertexed with $\text{rk } \Delta$ loops, and such an α is easily found.

Proof (of Proposition 9). That $p^{-1}(T)$ is a forest follows by spur-lifting. For (ii), p is injective on the vertices (and hence edges) of T_i as T is a tree and by spur lifting; surjectivity follows by path lifting. If q', q are the top and bottom quotient maps, define for any cell $q'(x) \in \Lambda/T_i$, the map $p'q'(x) = qp(x)$. Taking a vertex $v \in T$, the degree assertion follows immediately from (ii). \square

A covering $p : \tilde{\Delta}_u \rightarrow \Delta_v$ is *universal* iff for any covering $r : \Gamma_w \rightarrow \Delta_v$ there is a covering $q : \tilde{\Delta}_u \rightarrow \Gamma_w$ with $p = rq$. Equivalently, p is universal when any other covering of Δ is intermediate to it.

To construct a universal covering, one mimics a standard construction in topology, taking as the vertices the homotopy classes of paths in Δ starting at v . There is an edge \tilde{e} of $\tilde{\Delta}$ with start vertex the class of γ_1 and finish vertex the class of γ_2 if and only if there is an edge e of Δ with $\gamma_1 e$ homotopic to γ_2 . Define $\tilde{\Delta}_{[v]} \rightarrow \Delta_v$ by sending the class of γ to $t(\gamma)$ and the edge \tilde{e} described above to e .

Proposition 10. $\tilde{\Delta}_{[v]}$ is connected, simply connected, and the map $\tilde{\Delta}_{[v]} \rightarrow \Delta_v$ is a universal covering.

Proof. If $[\gamma]$ is a vertex of $\tilde{\Delta}_{[v]}$ with $\gamma = e_1 e_2 \dots e_k$, then $\tilde{e}_1 \tilde{e}_2 \dots \tilde{e}_k$ is a path from $[v]$ to $[\gamma]$ and so $\tilde{\Delta}_{[v]}$ is connected. That p is a covering is straight forward, and hence $\tilde{\Delta}_{[v]}$ is simply connected, for $\tilde{\gamma} = \tilde{e}_1 \tilde{e}_2 \dots \tilde{e}_k$ is a closed path at $[v]$ if and only if $e_1 e_2 \dots e_k$ is homotopic to v , ie: $\gamma = e_1 e_2 \dots e_k$ is homotopically trivial in Δ , giving $\tilde{\gamma}$ homotopically trivial as $p(\tilde{\gamma}) = \gamma$ and by homotopy lifting. If $r : \Gamma_w \rightarrow \Delta_v$ is a covering then define $q : \tilde{\Delta}_{[v]} \rightarrow \Gamma_w$ by $q[\gamma] = t(\gamma')$ where γ' is the lift via r of γ to the vertex w ; if \tilde{e} is an edge with start vertex $[\gamma]$ then let $q(\tilde{e})$ be the lift via r of $p(\tilde{e})$ to the vertex $t(\gamma')$. It is easy to see that $p = rq$ and hence q a covering by Lemma 3(i). \square

Many authors, anticipating the Galois correspondence, define a covering to be universal iff it is simply connected.

1.10. The lattice of intermediate coverings

Throughout this section Λ, Δ are connected graphs and $p : \Lambda_u \rightarrow \Delta_v$ is some fixed pointed covering. A connected pointed intermediate covering $\Lambda_u \rightarrow \Gamma_x \rightarrow \Delta_v$ is *equivalent* to another such, $\Lambda_u \rightarrow \Gamma_y \rightarrow \Delta_v$, if and only if there is a homeomorphism $\beta : \Gamma_x \rightarrow \Gamma_y$ such that

$$\begin{array}{ccccc} & & \Gamma_x & & \\ & \nearrow & \downarrow & \searrow & \\ \Lambda_u & & & & \Delta_v \\ & \searrow & \uparrow & \nearrow & \\ & & \Gamma_y & & \end{array}$$

commutes. Let $\mathcal{L}(\Lambda_u, \Delta_v)$ be the set of equivalence classes of such connected intermediate coverings.

Define $(\Lambda_u \rightarrow \Gamma_{x_1} \rightarrow \Delta_v) \leq (\Lambda_u \rightarrow \Upsilon_{x_2} \rightarrow \Delta_v)$, or just $\Gamma_{x_1} \leq \Upsilon_{x_2}$, if and only if there is a covering $s : \Upsilon_{x_2} \rightarrow \Gamma_{x_1}$ with $p = r_1 s q_2$, where r_1 is the covering $\Gamma_{x_1} \rightarrow \Delta_v$ and q_2 is $\Lambda_u \rightarrow \Upsilon_{x_2}$. If $\beta_1 : \Gamma_{x_1} \rightarrow \Psi_{y_1}$ and $\beta_2 : \Upsilon_{x_2} \rightarrow \Phi_{y_2}$ are homeomorphisms realizing equivalent coverings, then $\beta_2^{-1} s \beta_1 : \Phi_{y_2} \rightarrow \Psi_{y_1}$ is a covering with $p = r'_1 (\beta_2^{-1} s \beta_1) q'_2$. Thus, \leq is well defined upto equivalence, giving $\mathcal{L}(\Lambda_u, \Delta_v)$ the structure of a poset. We will also write $\Gamma_w \in \mathcal{L}(\Lambda_u, \Delta_v)$ for an equivalence class of intermediate coverings, without reference to the intermediate covering maps.

Recall that a poset (\mathcal{L}, \leq) equipped with a join \vee (or supremum) and meet \wedge (or infimum) is a lattice. A $\hat{0}$ (resp. $\hat{1}$) is an element such that $\hat{0} \leq x$ (resp. $x \leq \hat{1}$) for all x , and a lattice isomorphism (resp. anti-isomorphism) $\Theta : \mathcal{L}_1 \rightarrow \mathcal{L}_2$ is an order-preserving (resp. order-reversing) bijection whose inverse is also order-preserving (resp. order-reversing). In particular, an isomorphism sends joins to joins (and meets to meets) and an anti-isomorphism sends joins to meets (and meets to joins). A canonical example is the subgroups of a group G , ordered by inclusion, and with $A \vee B = \langle A, B \rangle$, $A \wedge B = A \cap B$, $\hat{0}$ the trivial subgroup and $\hat{1} = G$. The remainder of this section is devoted to showing that $\mathcal{L}(\Lambda_u, \Delta_v)$ is a lattice.

Let $\Lambda_u \rightarrow \Gamma_{x_1} \rightarrow \Delta_v$ and $\Lambda_u \rightarrow \Upsilon_{x_2} \rightarrow \Delta_v$ be intermediate to p , and q the quotient map used in the construction of the pushout of the coverings $q_1 : \Lambda_u \rightarrow \Gamma_{x_1}$ and $q_2 : \Lambda_u \rightarrow \Upsilon_{x_2}$. Let $x = q(x_1) = q(x_2)$ and $(\Gamma \amalg_{\Lambda} \Upsilon)_x$ the resulting pointed pushout.

Proposition 11. (i). *We have the intermediate covering $\Lambda_u \xrightarrow{t_i q_i} (\Gamma \amalg_{\Lambda} \Upsilon)_x \xrightarrow{r} \Delta_v$, where r is provided by the universality of the pushout.*

(ii). *Let $\Psi_{y_1}, \Phi_{y_2} \in \mathcal{L}(\Lambda_u, \Delta_v)$ be equivalent to $\Gamma_{x_1}, \Upsilon_{x_2}$ with $\beta_1 : \Gamma_{x_1} \rightarrow \Psi_{y_1}$, $\beta_2 : \Upsilon_{x_2} \rightarrow \Phi_{y_2}$ the corresponding homeomorphisms and $\beta_1 \amalg \beta_2 : \Gamma \amalg \Upsilon \rightarrow \Psi \amalg \Phi$ (disjoint unions) defined by $\beta_1 \amalg \beta_2|_{\Gamma} = \beta_1$ and $\beta_1 \amalg \beta_2|_{\Upsilon} = \beta_2$. Then the map $\beta : (\Gamma \amalg_{\Lambda} \Upsilon)_x \rightarrow (\Psi \amalg_{\Lambda} \Phi)_y$ defined by $\beta q = q'(\beta_1 \amalg \beta_2)$ is a homeomorphism making these pointed pushouts equivalent.*

Thus there is a well defined pushout of two elements of $\mathcal{L}(\Lambda_u, \Delta_v)$. As the proof will show, the maps $t_1, t_2 : \Gamma_{x_1}, \Upsilon_{x_2} \rightarrow (\Gamma \amalg_{\Lambda} \Upsilon)_x$ are coverings, and so the pushout is a lower bound for $\Gamma_{x_1}, \Upsilon_{x_2} \in (\mathcal{L}(\Lambda_u, \Delta_v), \leq)$, and the universality implies that it is an infimum.

Proof (of Proposition 11). (i). If $v, v' \in \Gamma_1 \amalg \Gamma_2$ are vertices with $v \sim v'$ and e an edge with start v then, by successively lifting and covering, one can show that there is an edge e' with start v' such that $e \sim e'$. Now, v_1 maps via t_1 to $[v_1]$ and if $[e']$ an edge starting at this vertex then $s(e') \sim v_1$, and so by the above there is an edge e starting at v_1 with $[e'] = [e]$. Thus $t_1(e) = [e']$, and so t_1 maps the edges starting at v_1 surjectively onto those starting at $t_1(v_1)$. If e, e' are edges starting at v_1 with $t_1(e) = t_1(e')$, then one gets by induction that $q_1(e) = q_1(e')$, and q_1 a covering forces $e = e'$, and thus t_1 (similarly t_2) is a covering, hence the $t_i q_i : \Lambda_u \rightarrow (\Gamma \amalg_{\Lambda} \Upsilon)_x$ are too. The map r is provided by the universality and is a covering by Lemma 3(i). Part (ii) is a tedious but routine diagram chase. \square

Now to pullbacks. With $\Lambda_u \rightarrow \Gamma_{x_1} \rightarrow \Delta_v$ and $\Lambda_u \rightarrow \Upsilon_{x_2} \rightarrow \Delta_v$ intermediate to p , $x = x_1 \times x_2$ is a vertex of the pullback of the coverings $r_1 : \Gamma_{x_1} \rightarrow \Delta_v$ and $r_2 : \Upsilon_{x_2} \rightarrow \Delta_v$. Let $(\Gamma \amalg_{\Delta} \Upsilon)_x$ be the pointed pullback consisting of the component containing the vertex x .

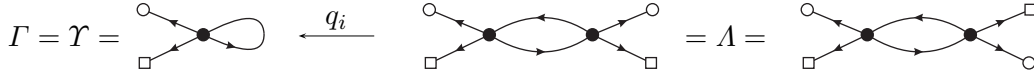
Proposition 12. (i). *We have the intermediate covering $\Lambda_u \xrightarrow{q} (\Gamma \amalg_{\Delta} \Upsilon)_x \xrightarrow{r_i t_i} \Delta_v$, where q is provided by the universality of the pullback.*

(ii). *Let $\Psi_{y_1}, \Phi_{y_2} \in \mathcal{L}(\Lambda_u, \Delta_v)$ be equivalent to $\Gamma_{x_1}, \Upsilon_{x_2}$ with β_1, β_2 the corresponding homeomorphisms and $\beta : (\Gamma \amalg_{\Delta} \Upsilon)_x \rightarrow (\Psi \amalg_{\Delta} \Phi)_y$ defined by $\beta(x \times y) = \beta_1(x) \times \beta_2(y)$. Then β is a homeomorphism making the pointed pullbacks equivalent.*

Thus there is a well defined pullback of two elements of $\mathcal{L}(\Lambda_u, \Delta_v)$. Again the proof shows that the maps $t_1, t_2 : (\Gamma \amalg_{\Delta} \Upsilon)_x \rightarrow \Gamma_{x_1}, \Upsilon_{x_2}$ are coverings, and so the pullback is an upper bound for $\Gamma_{x_1}, \Upsilon_{x_2} \in (\mathcal{L}(\Lambda_u, \Delta_v), \leq)$, and the universality implies that it is a supremum.

Proof (of Proposition 12). (i). We show that t_1 is a covering; t_2 is similar. From $t_1(e_1 \times e_2) = e_1$ it is clear that t_1 is dimension preserving. For $t_1(u_1 \times u_2) = u_1$, let e_1 be an edge of Γ with $s(e_1) = u_1$. Then e_1 covers $r_1(e_1)$ which lifts via the covering r_2 to u_2 in \mathcal{Y} to an edge e_2 covering $r_1(e_1)$, ie: with $r_2(e_2) = r_1(e_1)$. Thus there is an edge $e_1 \times e_2$ of the pullback with $s(e_1 \times e_2) = u_1 \times u_2$ and $t_1(e_1 \times e_2) = e_1$, giving the surjectivity of t_1 on the edges starting at u_1 . If $e'_1 \times e'_2$ starts at $u_1 \times u_2$ and $t_1(e'_1 \times e'_2) = e_1$ then $e'_1 = e_1$. We have $t_2(e'_1 \times e'_2) = e'_2$ starting at u_2 , and $r_2(e'_2) = r_1 t_1(e'_1 \times e'_2) = r_1(e_1) = r_2(e_2)$. Thus, as r_2 is a cover, we have $e'_2 = e_2$ and so $e'_1 \times e'_2 = e_1 \times e_2$, and t_1 is indeed a covering. Hence the $r_i t_i$ are too and q by Lemma 3(i). Part (ii) is an analogous to that of Proposition 11. \square

The proof of Proposition 12 also shows that the $t_1, t_2 : \Gamma \amalg_{\Delta} \mathcal{Y} \rightarrow \Gamma, \mathcal{Y}$ are coverings in the unpointed case. We pause to observe a slight asymmetry to the duality between pushouts and pullbacks: given coverings $r_1, r_2 : \Gamma, \mathcal{Y} \rightarrow \Delta$, the $t_1, t_2 : \Gamma \amalg_{\Delta} \mathcal{Y} \rightarrow \Gamma, \mathcal{Y}$ are coverings, whereas coverings $q_1, q_2 : \Lambda \rightarrow \Gamma, \mathcal{Y}$ do not necessarily give coverings $t_1, t_2 : \Gamma, \mathcal{Y} \rightarrow \Gamma \amalg_{\Lambda} \mathcal{Y}$, unless the q_i are intermediate $\Lambda \rightarrow (\Gamma \text{ or } \mathcal{Y}) \rightarrow \Delta$. Indeed, taking the $\Gamma = \mathcal{Y}$ to be two copies of the left hand graph,



and the coverings $q_i : \Lambda \rightarrow \Gamma$ or \mathcal{Y} (described here by drawing the fibers of the vertices), then the t_i provided by the pushout construction are not coverings of the pushout.

Summarising the results of this section:

Theorem 1 (lattice of intermediate coverings). $\mathcal{L}(\Lambda_u, \Delta_v)$ is a lattice with join $\Gamma_{x_1} \vee \mathcal{Y}_{x_2}$ the pullback $(\Gamma \amalg_{\Delta} \mathcal{Y})_{x_1 \times x_2}$, meet $\Gamma_{x_1} \wedge \mathcal{Y}_{x_2}$ the pushout $(\Gamma \amalg_{\Lambda} \mathcal{Y})_{q(x_i)}$, $\widehat{0} = \Delta_v$ and $\widehat{1} = \Lambda_u$.

The pointing of the covers in this section is essential if one wishes to work with *connected* intermediate coverings and also have a lattice structure (both of which we do). The problem is the pullback: because it is not in general connected, we need the pointing to tell us which component to choose.

2. The Galois theory of graphs

The ‘‘Galois correspondence’’ between coverings of graphs and subgroups of the fundamental group goes back to Reidemeister [16] (see eg: [3]). We provide a slightly alternative formulation that exploits the lattice structure of §1.10 and is more in the spirit of classical Galois theory.

Throughout this section $p : \Lambda_u \rightarrow \Delta_v$ is a fixed covering with Λ, Δ connected. An *automorphism* (or *deck transformation*) of p is a graph homeomorphism $\alpha : \Lambda_u \rightarrow \Lambda_u$ making the diagram,

$$\begin{array}{ccc} \Lambda_u & \xrightarrow{\alpha} & \Lambda_u \\ & \searrow p & \swarrow p \\ & \Delta_v & \end{array}$$

commute. The automorphisms form a group $\text{Gal}(\Lambda_u, \Delta_v) = \text{Gal}(\Lambda_u \xrightarrow{p} \Delta_v)$, the *Galois group* of the covering.

Lemma 4. (i). The action of $\text{Gal}(\Lambda_u, \Delta_v)$ on Λ is orientation preserving. (ii). The effect of an automorphism $\alpha : \Lambda_u \rightarrow \Lambda_u$ is completely determined by $\alpha(u) = u'$. In particular, the Galois group acts freely on Λ .

Proof. (i). Both the edge e and $\alpha(e)$ lie in the same fiber of the covering, so that if $\alpha(e) = e^{-1}$ then $p(e) = p(e)^{-1}$, a contradiction, so the Galois group acts without inversions. (ii). If x is a vertex of Λ and γ a path from u to x , then $\alpha(x)$ is the terminal vertex of the lift to u' of the path $p(\gamma)$. The images of the edges are handled similarly. \square

The explicit construction of automorphisms is achieved by the following technical result:

Proposition 13. *Let $p : \Lambda_u \rightarrow \Delta_v$ be a covering and u' another vertex in the fiber of v such that for any closed path γ at v with lifts γ_i at u, u' , we have γ_1 closed if and only if γ_2 closed. For any vertex $x \in \Lambda$ and path μ from u to x , let $\alpha(x)$ be the terminal vertex of the lift at u' of $p(\mu)$. Then $x \mapsto \alpha(x)$ extends to an automorphism $\alpha \in \text{Gal}(\Lambda_u, \Delta_v)$.*

In particular, for a covering satisfying (i) and (ii) of the Proposition, there is an element of the Galois group sending the vertex u to the vertex u' .

Proof. If μ' is another path from u to x , then $p(\mu)p(\mu')^{-1}$ is a closed path at v that lifts to a closed path (ie: $\mu(\mu')^{-1}$) at u , hence to a closed path at u' . Thus $\alpha(x)$ is also the terminal vertex of the lift at u' of $p(\mu')$ and α is a well defined map $V_\Lambda \rightarrow V_\Lambda$. To extend α to the edges, let μ be a path from u to the vertex $s(e)$, and lift the path $p(\mu e)$ to u' . Define $\alpha(e)$ to be the lift of $p(e)$ to the terminal vertex of $p(\mu)$. It is easy to see that $\alpha : \Lambda \rightarrow \Lambda$ is a surjective dimension-preserving map of graphs, and that x and $\alpha(x)$ lie in the same fiber of the covering, for any cell x , whence $p\alpha = p$.

It remains to show that α is injective. For vertices x and x' , choose paths μ, μ' from u to x and x' . Then if $\alpha(x) = \alpha(x')$, the lifts at u' of $p(\mu)$ and $p(\mu')$ finish at the same vertex, and so therefore must $p(\mu)$ and $p(\mu')$, as p is well defined at the vertex $\alpha(x) = \alpha(x')$. Thus μ, μ' finish at the same vertex and so $x = x'$. For edges e, e' with $\alpha(e) = \alpha(e')$, the injectivity of α on the vertices gives that they must have the same start vertex, and moreover must lie in the same fiber, hence $e = e'$, by the injectivity of coverings on the edges with start a given vertex. \square

Let $\Lambda_u \rightarrow \Gamma_x \rightarrow \Delta_v$ be a covering intermediate to p and consider those $\alpha \in \text{Gal}(\Lambda_u, \Delta_v)$ such that

$$\begin{array}{ccc} \Lambda_u & \xrightarrow{\alpha} & \Lambda_{u'} \\ & \searrow q & \swarrow q \\ & \Gamma_x & \end{array}$$

commutes. This gives a subgroup that can be identified with $\text{Gal}(\Lambda_u \xrightarrow{q} \Gamma_x)$. If $\Lambda_u \rightarrow \Upsilon_y \rightarrow \Delta_v$ is an equivalent covering with homeomorphism $\beta : \Gamma_x \rightarrow \Upsilon_y$, then

$$\begin{array}{ccc} \Lambda_u & \xrightarrow{\alpha} & \Lambda_{u'} \\ & \searrow q & \swarrow q \\ & \Gamma_x & \end{array} \text{ commutes} \Leftrightarrow \begin{array}{ccc} \Lambda_u & \xrightarrow{\alpha} & \Lambda_{u'} \\ & \searrow q' & \swarrow q' \\ & \Upsilon_y & \end{array} \text{ commutes.}$$

Thus $\text{Gal}(\Lambda_u \xrightarrow{q} \Gamma_x) = \text{Gal}(\Lambda_u \xrightarrow{q'} \Upsilon_y)$, and we can associate in a well defined manner a subgroup of the Galois group to an element of the lattice $\mathcal{L}(\Lambda_u, \Delta_v)$.

On the other hand, if $H \subset \text{Gal}(\Lambda_u, \Delta_v)$, then by Lemma 4 we may form the quotient Λ/H , and indeed,

Lemma 5. *If $H_1 \subset H_2 \subset \text{Gal}(\Lambda_u, \Delta_v)$ then,*

$$\Lambda_u \xrightarrow{q_1} (\Lambda/H_1)_{q_1(u)} \xrightarrow{s} (\Lambda/H_2)_{q_2(u)} \xrightarrow{r} \Delta_v$$

are all coverings, where $q_i : \Lambda \rightarrow \Lambda/H_i$ are the quotient maps and s, r are defined by $sq_1 = q_2$ and $rq_2 = p$.

Proof. All the complexes are connected, and the q_i coverings by Lemmas 4 and 3(ii). Thus s is a covering by Lemma 3(i), and another application, this time to $p = r(sq_1)$, gives that r is a covering. \square

Thus, letting $H_1 = H_2 = H$, we can associate to $H \subset \text{Gal}(\Lambda_u, \Delta_v)$ the intermediate covering $\Lambda_u \rightarrow (\Lambda/H)_{q(u)} \rightarrow \Delta_v$, and by passing to its equivalence class, we get an element of the lattice $\mathcal{L}(\Lambda_u, \Delta_v)$ associated to H .

Proposition 14. *The following are equivalent for a covering $p : \Lambda_u \rightarrow \Delta_v$:*

1. *For all closed paths γ at v , the lifts of γ to each vertex of $\text{fib}_{\Lambda \rightarrow \Delta}(v)$ are either all closed or all non-closed;*
2. *$\text{Gal}(\Lambda_u \xrightarrow{p} \Delta_v)$ acts regularly on $\text{fib}_{\Lambda \rightarrow \Delta}(v)$.*

In any case, we call the covering $\Lambda_u \rightarrow \Delta_v$ *Galois*, with *regular* a common alternative as the second part of the Proposition makes clear. It is clear that if $\Lambda_u \rightarrow \Delta_v$ is Galois then so is $\Lambda_{u'} \rightarrow \Delta_v$ for any other u' in the fiber of v ; if $\Lambda_u \rightarrow \Gamma_w \rightarrow \Delta_v$ is intermediate with $\Lambda_u \rightarrow \Delta_v$ Galois, then $\Lambda_u \rightarrow \Gamma_w$ is Galois.

Proof. The equivalence follows immediately from Proposition 13 and the fact that automorphisms send closed paths to closed paths and non-closed paths to non-closed paths. \square

If the covering $p : \Lambda_u \rightarrow \Delta_v$ is Galois, let $g \in \pi_1(\Delta, v)$ with representative path γ and α_g an automorphism in $\text{Gal}(\Lambda_u, \Delta_v)$ that sends u to the terminal vertex u' of the lift of γ to u . By homotopy lifting, α_g depends only on the vertices u, u' and not on the chosen representative path γ , and so the map $\pi_1(\Delta, v) \rightarrow \text{Gal}(\Lambda_u, \Delta_v)$ given by $g \mapsto \alpha_g$ is well defined.

Proposition 15. *If $p : \Lambda_u \rightarrow \Delta_v$ is Galois then $g \mapsto \alpha_g$ is a surjective homomorphism with kernel $p^*\pi_1(\Lambda, u)$, such that under the induced isomorphism*

$$\pi_1(\Delta, v)/p^*\pi_1(\Lambda, u) \rightarrow \text{Gal}(\Lambda_u, \Delta_v),$$

if $\Lambda_u \xrightarrow{q} \Gamma_x \xrightarrow{r} \Delta_v$ is intermediate, then the subgroup $r^\pi_1(\Gamma, x)/p^*\pi_1(\Lambda, u)$ has image $\text{Gal}(\Lambda_u \xrightarrow{q} \Gamma_x)$.*

Proof. It is easy to check that $\alpha_{g_1 g_2} = \alpha_{g_1} \alpha_{g_2}$ and so we have a homomorphism. If $\beta \in \text{Gal}(\Lambda_u, \Delta_v)$ then β is the unique automorphism sending u to $\beta(u)$. Taking a path μ in Λ from u to $\beta(u)$ thus gives $\beta = \alpha_g$ for g the homotopy class of $p(\mu)$, and hence the homomorphism is surjective. Because automorphisms act freely, an element g is in the kernel iff α_g fixes the vertex u , and this happens precisely when g can be represented by a path lifting to a closed path at u , ie: when $g \in p^*\pi_1(\Lambda, u)$. It is easy to check that this homomorphism maps $r^*\pi_1(\Gamma, x)$ onto $\text{Gal}(\Lambda_u, \Gamma_x)$. \square

Corollary 1. *A covering $\Lambda_u \xrightarrow{p} \Delta_v$ is Galois if and only if $p^*\pi_1(\Lambda, u)$ is a normal subgroup of $\pi_1(\Delta, v)$.*

Proof. It remains to show the “if” part. Let u' be a vertex in the fiber of v , γ a closed path at v with lifts γ_1, γ_2 at u and u' , and μ a path from u to u' . Let g and h be the homotopy classes of γ and $p(\mu)$. Then γ_1 is closed iff $g \in p^*\pi_1(\Lambda, u) \Leftrightarrow hgh^{-1} \in p^*\pi_1(\Lambda, u)$ by normality, and this in turn happens precisely when $\mu\gamma_1\mu^{-1}$ is closed at u , ie: when γ_2 is closed at u' . Thus the covering is Galois. \square

Proposition 16. *Let $p : \Lambda_u \rightarrow \Delta_v$ be a Galois covering. If $H \subset \text{Gal}(\Lambda_u, \Delta_v)$ then,*

$$[\text{Gal}(\Lambda_u, \Delta_v) : H] = \deg(\Lambda/H_{q(u)} \xrightarrow{r} \Delta_v).$$

Proof. If a group G acts regularly on a set and H is a subgroup, then the number of H -orbits is the index $[G : H]$. The result follows as the H -orbits on the fiber (via p) of v are precisely the vertices of Λ/H covering v (via r). \square

In particular the Galois group of a Galois covering has order the degree of the covering. We have now assembled sufficient machinery to prove,

Theorem 2 (Galois correspondence). *Let $\Lambda_u \rightarrow \Delta_v$ be a Galois covering with $\mathcal{L}(\Lambda_u, \Delta_v)$ the lattice of equivalence classes of intermediate coverings and $\text{Gal}(\Lambda_u, \Delta_v)$ the Galois group. Then the map that associates to $\Lambda_u \rightarrow \Gamma_x \rightarrow \Delta_v \in \mathcal{L}(\Lambda_u, \Delta_v)$ the subgroup $\text{Gal}(\Lambda_u, \Gamma_x)$ is a lattice anti-isomorphism from $\mathcal{L}(\Lambda_u, \Delta_v)$ to the lattice of subgroups of $\text{Gal}(\Lambda_u, \Delta_v)$. Its inverse is the map associating to $H \subset \text{Gal}(\Lambda_u, \Delta_v)$ the element $\Lambda_u \rightarrow \Lambda/H_{q(u)} \rightarrow \Delta_v \in \mathcal{L}(\Lambda_u, \Delta_v)$.*

Proof. Let f and g be the two maps described in the theorem. It is easiest to work from the point of view of g : if $H_1 \leq H_2$ in the lattice of subgroups, then the covering s of Lemma 5 gives $g(H_2) \leq g(H_1)$, so g is an anti-morphism of lattices. If $\Lambda_u \rightarrow \Gamma_x \rightarrow \Delta_v$ is intermediate, then we also have the intermediate covering $\Lambda_u \rightarrow \Lambda_u/\text{Gal}(\Lambda_u, \Gamma_x) \rightarrow \Gamma_x$ with $\Lambda_u \rightarrow \Gamma_x$ is Galois. By Proposition 16, the covering $\Lambda_u/\text{Gal}(\Lambda_u, \Gamma_x) \rightarrow \Gamma_x$ has degree 1, hence is a homeomorphism (Lemma 2(iii)) and we have the diagram at right, with the whole square and the left triangle commuting by intermediacy, hence the right triangle commuting as well. Thus, the intermediate coverings $\Lambda_u \rightarrow \Gamma_x \rightarrow \Delta_v$ and $\Lambda_u \rightarrow \Lambda_u/\text{Gal}(\Lambda_u, \Gamma_x) \rightarrow \Delta_v$ are equivalent, and we have $gf = \text{id}$. If $H \subset \text{Gal}(\Lambda_u, \Delta_v)$ and $q : \Lambda \rightarrow \Lambda/H$ the quotient map, then $q\alpha = q$ for any $\alpha \in H$ and so $H \subset \text{Gal}(\Lambda_u, (\Lambda/H)_{q(u)})$ with the covering $\Lambda_u \rightarrow \Lambda/H_{q(u)}$ intermediate, hence Galois. Proposition 16 gives the index of H in $\text{Gal}(\Lambda_u, (\Lambda/H)_{q(u)})$ to be the degree of the covering $\Lambda/H_{q(u)} \rightarrow \Lambda/H_{q(u)}$, ie: $H = \text{Gal}(\Lambda_u, (\Lambda/H)_{q(u)})$, and we have $fg = \text{id}$. \square

$$\begin{array}{ccccc} & & \Lambda/\text{Gal}(\Lambda_u, \Gamma_x) & & \\ & \nearrow & \downarrow & \searrow & \\ \Lambda_u & & \Gamma_x & & \Delta_v \end{array}$$

As lattice anti-isomorphisms send joins to meets and meets to joins we have as an immediate corollary that,

Corollary 2. *Let $\Lambda_u \rightarrow \Delta_v$ be Galois with $\Lambda_u \rightarrow \Gamma_x \rightarrow \Delta_v$ and $\Lambda_u \rightarrow \Upsilon_y \rightarrow \Delta_v$ in the lattice $\mathcal{L}(\Lambda_u, \Delta_v)$ and $H_1, H_2 \subset \text{Gal}(\Lambda_u, \Delta_v)$. Then,*

$$\begin{aligned} \text{Gal}(\Lambda_u, (\Gamma \prod_{\Delta} \Upsilon)_z) &= \text{Gal}(\Lambda_u, \Gamma_x) \cap \text{Gal}(\Lambda_u, \Upsilon_y), \\ \text{Gal}(\Lambda_u, (\Gamma \prod_{\Lambda} \Upsilon)_z) &= \langle \text{Gal}(\Lambda_u, \Gamma_x), \text{Gal}(\Lambda_u, \Upsilon_y) \rangle, \end{aligned}$$

and the intermediate coverings,

$$\begin{aligned} \Lambda_u \rightarrow \Lambda/\langle H_1, H_2 \rangle_{t(u)} \rightarrow \Delta_v \text{ and } \Lambda_u \rightarrow (\Lambda/H_1 \prod_{\Lambda} \Lambda/H_2)_{qq_i(u)} \rightarrow \Delta_v, \\ \Lambda_u \rightarrow \Lambda/(H_1 \cap H_2)_{t(u)} \rightarrow \Delta_v \text{ and } \Lambda_u \rightarrow (\Lambda/H_1 \prod_{\Delta} \Lambda/H_2)_w \rightarrow \Delta_v. \end{aligned}$$

are equivalent (where $t : \Lambda \rightarrow \Lambda/\langle H_1, H_2 \rangle$ or $\Lambda/(H_1 \cap H_2)$, $q_i : \Lambda \rightarrow \Lambda/H_i$ and q the quotient from the pushout).

(This result is essentially Theorems 4.3 and 5.5 of [20], restated in our terms.)

The universal cover $\tilde{\Delta}_u \rightarrow \Delta_v$ ($u = [v]$) is Galois by Proposition 10 and Corollary 1, and by Proposition 15 there is an isomorphism

$$\varphi : \pi_1(\Delta, v) \rightarrow \text{Gal}(\tilde{\Delta}_u, \Delta_v)$$

such that if $\tilde{\Delta}_u \rightarrow \Gamma_x \xrightarrow{r} \Delta_v$ is intermediate, then the subgroup $r^*\pi_1(\Gamma, x)$ has image $\text{Gal}(\tilde{\Delta}_u, \Gamma_x)$. Moreover, two intermediate coverings,

$$\tilde{\Delta}_u \rightarrow \Gamma_x \xrightarrow{r} \Delta_v \text{ and } \tilde{\Delta}_u \rightarrow \Upsilon_y \xrightarrow{r'} \Delta_v,$$

are equivalent if and only if there is a homeomorphism $\beta : \Gamma_x \rightarrow \Upsilon_y$ with $r = r'\beta$. We thus obtain the more familiar version of the Galois correspondence, as a special case of Theorem 2:

Corollary 3 (Galois correspondence for the universal cover). *The map that associates to a covering $r : \Gamma_x \rightarrow \Delta_v$ the subgroup $r^*\pi_1(\Gamma, x)$ is a lattice anti-isomorphism from $\mathcal{L}(\tilde{\Delta}_u, \Delta_v)$ to the lattice of subgroups of $\pi_1(\Delta, v)$ that sends Galois covers to normal subgroups. Its inverse associates to $H \subset \pi_1(\Delta, v)$ the covering $\tilde{\Delta}/\varphi(H)_{q(u)} \rightarrow \Delta_v$.*

We end this section by showing that the excision of trees has little effect on the lattice $\mathcal{L}(\Lambda, \Delta)$. Let $p : \Lambda_u \rightarrow \Delta_v$ be a covering, $T \hookrightarrow \Delta$ a spanning tree, $T_i \hookrightarrow \Lambda$ the components of $p^{-1}(T)$ and $p : (\Lambda/T_i)_{q(u)} \rightarrow (\Delta/T)_{q(v)}$ the induced covering (where we have (ab)used q for both quotients and p for both coverings).

Theorem 3 (lattice excision). *There is a degree and rank preserving isomorphism of lattices*

$$\mathcal{L}(\Lambda, \Delta) \rightarrow \mathcal{L}(\Lambda/T_i, \Delta/T),$$

that sends the equivalence class of $\Lambda_u \rightarrow \Gamma_x \xrightarrow{r} \Delta_v$ to the equivalence class of $\Lambda/T_i \rightarrow \Gamma/T'_i \rightarrow \Delta/T$ (with $T'_i \hookrightarrow \Gamma$ the components of $r^{-1}(T)$) and Galois coverings to Galois coverings.

This result could have been shown directly and messily at the end of §1.10; we use the Galois correspondence.

Proof. The quotient $q : \Delta \rightarrow \Delta/T$ induces an isomorphism $q^* : \pi_1(\Delta, v) \rightarrow \pi_1(\Delta/T, q(v))$ where $(qp)^* = (pq)^*$ by the commutativity of the diagram in Proposition 9. Thus $q^*p^*\pi_1(\Lambda, v) = p^*\pi_1(\Lambda/T_i, q(u))$ giving that $\Lambda \rightarrow \Delta$ is Galois iff $\Lambda/T_i \rightarrow \Delta/T$ is Galois by Corollary 1, and an isomorphism $\text{Gal}(\Lambda, \Delta) \rightarrow \text{Gal}(\Lambda/T_i, \Delta/T)$ by Proposition 15 (leaving off the pointings for clarity). This in turn induces an isomorphism $\mathcal{L}_1 \rightarrow \mathcal{L}_2$ between the subgroup lattices of these two groups, so that two applications of the Galois correspondence gives

$$\mathcal{L}(\Lambda, \Delta) \rightarrow \mathcal{L}_1 \rightarrow \mathcal{L}_2 \rightarrow \mathcal{L}(\Lambda/T_i, \Delta/T),$$

a composition of an isomorphism and two anti-isomorphisms, hence the result in the case that $\Lambda \rightarrow \Delta$ is Galois.

If $\Lambda_u \rightarrow \Gamma_x \xrightarrow{r} \Delta_v$ is intermediate then q^* sends the subgroup $r^*\pi_1(\Gamma, x)$ to $r^*\pi_1(\Gamma/T'_i, q(x))$ and so the isomorphism of Galois groups sends $\text{Gal}(\Lambda, \Gamma)$ to $\text{Gal}(\Lambda/T_i, \Gamma/T'_i)$. This gives the desired image of the intermediate covering, but also, if $\Lambda \rightarrow \Delta$ is not Galois, then $\mathcal{L}(\Lambda, \Delta)$ embeds as a sublattice of $\mathcal{L}(\tilde{\Delta}, \Delta)$, sent to the sublattice $\mathcal{L}(\Lambda/T_i, \Delta/T)$ via the result applied to the Galois covering $\tilde{\Delta} \rightarrow \Delta$. \square

Thus in particular, there are homeomorphisms

$$\begin{aligned} (\Gamma \coprod_{\Delta} \Upsilon)/T_k &\rightarrow (\Gamma/T_{1j}) \coprod_{\Delta/T} (\Upsilon/T_{2j}), \\ (\Gamma \coprod_{\Lambda} \Upsilon)/T_k &\rightarrow (\Gamma/T_{1j}) \coprod_{\Lambda/T_i} (\Upsilon/T_{2j}). \end{aligned}$$

where Γ, Υ are intermediate to $\Lambda \rightarrow \Delta$, and the trees T_{ij}, T_{2j}, T_k are the components of the preimages of T via the various coverings.

3. Graphs of finite rank

This section is devoted to a more detailed study of the form of those covering graphs $\Lambda \rightarrow \Delta$ where $\text{rk } \Lambda < \infty$.

Lemma 6. *Let Γ be a connected graph. (i). If $\text{rk } \Gamma < \infty$ and Θ the trivial graph, then $\text{rk } \bigvee_{\Theta} \Gamma = \text{rk } \Gamma + 1$. (ii). If Γ_1, Γ_2 are connected of finite rank, and Θ finite, then*

$$\text{rk} \left(\Gamma_1 \bigvee_{\Theta} \Gamma_2 \right) = |V_{\Theta}| - 1 + \sum \text{rk } \Gamma_i.$$

Proof. Part (i) follows from the comments at the end of §1.7, and (ii) by induction on $|\Theta|$ and (i). \square

Proposition 17. *A connected graph Λ has finite rank if and only if Λ decomposes as a wedge sum $\Lambda = \Gamma \bigvee_{\Theta} \Phi$ with Γ finite, locally finite, connected, Θ finite, Φ a forest, and no two vertices of the image of $\Theta \hookrightarrow \Phi$ lying in the same component.*

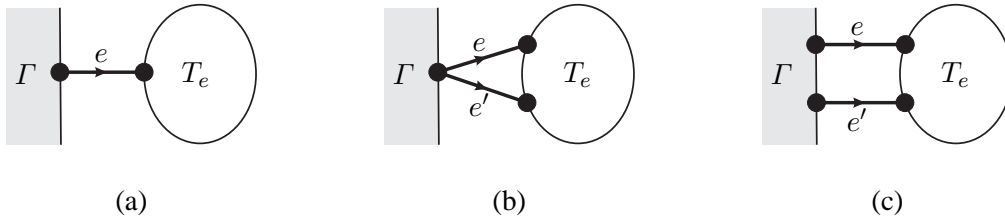
Proof. If Λ has such a decomposition then Φ necessarily has finitely many components and finite rankness follows by inductively applying Lemma 6. For the converse, fix a basepoint vertex v and spanning tree T so that there is a finite set P of arcs of Λ not in T . Take paths in T from v to the start and terminal vertices of the arcs of P . Let Γ be the union of the edges in P and these paths and let Φ the result of removing from T the edges in the paths. \square

Lemma 7. *Λ connected is of finite rank if and only if for any vertex v , the spine $\hat{\Lambda}_v$ is finite, locally finite.*

Proof. If Λ has finite rank then we have the wedge sum decomposition of Proposition 17. If e is an edge not contained in Γ and γ a closed path at v containing e , then e is contained in a tree component T of Φ . As this component is wedged onto Γ at a single vertex, that part of γ contained in T is a closed path, hence contains a spur. Thus e is contained in no closed reduced path at v and the so spine $\hat{\Lambda}$ is a subgraph of Γ , hence finite, locally finite. Conversely, a finite spine has finite rank, hence so does Λ by Lemma 1(i). \square

Proposition 18. *Let Λ be a connected graph, $\Gamma \hookrightarrow \Lambda$ a connected subgraph and $v \in \Gamma$ a vertex such that every closed reduced path at v in Λ is contained in Γ . Then Λ has a wedge sum decomposition $\Lambda = \Gamma \bigvee_{\Theta} \Phi$ with Φ a forest and no two vertices of the image of $\Theta \hookrightarrow \Phi$ lying in the same component.*

Proof. Consider an edge e of $\Lambda \setminus \Gamma$ having at least one of its end vertices $s(e)$ or $t(e)$, in Γ . For definiteness we can assume, by relabeling the edges in the arc \bar{e} , that it is $s(e)$ that is a vertex of Γ . If $t(e) \in \Gamma$ then by traversing a reduced path in Γ from v to $s(e)$, crossing e and a reduced path in Γ from $t(e)$ to v , we get a closed reduced path not contained in Γ , a contradiction. Thus $t(e) \notin \Gamma$. Let T_e be the union of all the reduced paths in $\Lambda \setminus \{e\}$ starting at $t(e)$, so we have the situation as in (a):



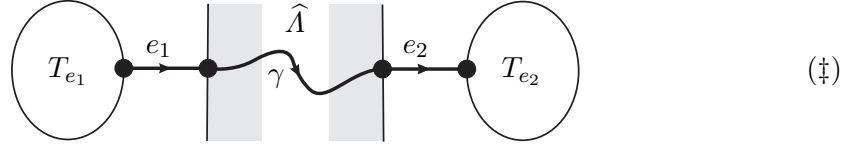
If γ is a non-trivial closed path in T_e starting at $t(e)$, then a path from v to $t(e)$, traversing γ , and going the same way back to v cannot be reduced. But the only place a spur can occur is in γ and so T_e is a tree. If e' is another edge of $\Lambda \setminus \Gamma$ with $s(e') \in \Gamma$ then we claim that neither of the two situations (b) and (c) above can occur, ie: $t(e')$ is not a vertex of T_e . For otherwise, a reduced closed path in T_e from $t(e)$ to $t(e')$ will give a reduced closed path at v not in Γ . Thus, another edge e' yields a tree $T_{e'}$ defined like T_e , but disjoint from it. Each component of Φ is thus obtained this way. \square

In particular we have such a decomposition involving a spine, and so Λ is made up of its spine at some vertex, together with a collection of trees, each connected to $\hat{\Lambda}_v$ by a single edge. If Λ has finite rank then $\hat{\Lambda}$ and Θ are finite, and we have

$$\Lambda = \left(\cdots \left(\left(\hat{\Lambda}_v \bigvee_{\Theta_1} T_1 \right) \bigvee_{\Theta_2} T_2 \right) \cdots \right) \bigvee_{\Theta_k} T_k, \quad (1)$$

with the Θ_i single vertices, the $\Theta_i \hookrightarrow \hat{\Lambda}_v$, and the images $\Theta_i \hookrightarrow T_i$ having valency one. Moreover, if $\Lambda \rightarrow \Delta$ is a covering with Δ single vertexed and Λ of finite rank, then by Proposition 8(i), each tree T_i realizes an embedding $\mathcal{R} \hookrightarrow \Lambda$ of the real line in Λ , and as the spine is finite,

the trees are thus paired



with the e_i (and indeed all the edges in the path $\mathcal{R} \hookrightarrow \Lambda$) in the same fiber of the covering. This pairing will play an important role in §4.

Corollary 4. *Let $\Lambda \rightarrow \Delta$ be a covering with Δ non-trivial, single vertexed and $\text{rk } \Lambda < \infty$. Then $\deg(\Lambda \rightarrow \Delta) < \infty$ if and only if $\Lambda = \hat{\Lambda}_v$.*

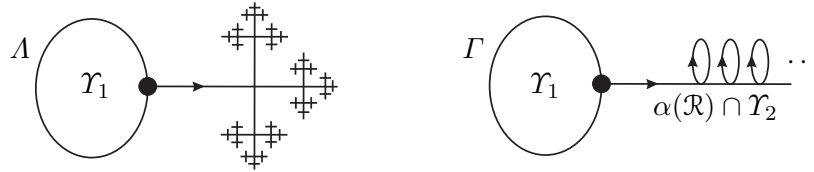
Proof. If Λ is more than $\hat{\Lambda}_v$ then one of the trees T_i in the decomposition (1) is non trivial and by Proposition 8(i) we get the real line $\mathcal{R} \hookrightarrow \Lambda$, with image in the fiber of an edge, contradicting the finiteness of the degree. The converse follows from Lemma 7. \square

Proposition 19. *Let $\Lambda \rightarrow \Delta$ be a covering with (i). $\text{rk } \Delta > 1$, (ii). $\text{rk } \Lambda < \infty$, and (iii). for any intermediate covering $\Lambda \rightarrow \Gamma \rightarrow \Delta$ we have $\text{rk } \Gamma < \infty$. Then $\deg(\Lambda \rightarrow \Delta) < \infty$.*

The covering $\mathcal{R} \rightarrow \Delta$ of a single vertexed Δ of rank 1 by the real line shows why the $\text{rk } \Delta > 1$ condition cannot be dropped.

Proof. By lattice excision, Theorem 3, we may pass to the Δ single vertexed case while preserving (i)-(iii). Establishing the degree here and passing back to the general Δ will give the result. If the degree of the covering $\Lambda \rightarrow \Delta$ is infinite for Δ single vertexed, then by Corollary 4, in the decomposition (1) for Λ , one of the trees is non-empty and Λ has the form of the graph in Proposition 8 with this non-empty tree the union of the edge e and \mathcal{Y}_2 .

Let Γ be a graph as defined as follows: take the union of \mathcal{Y}_1 , the edge e and $\alpha(\mathcal{R}) \cap \mathcal{Y}_2$, where $\alpha(\mathcal{R})$ is the embedding of the real line given by Proposition 8(i). At each vertex of $\alpha(\mathcal{R}) \cap \mathcal{Y}_2$ place $\text{rk } \Delta - 1$ edge loops:



(this picture depicting the $\text{rk } \Delta = 2$ case). Then there is an obvious covering $\Gamma \rightarrow \Delta$ so that by Proposition 8(ii) we have an intermediate covering $\Lambda \rightarrow \Gamma \rightarrow \Delta$. Equally obviously, Γ has infinite rank, contradicting (iii). Thus, $\deg(\Lambda \rightarrow \Delta) < \infty$. \square

Proposition 20. *Let $\Psi \rightarrow \Lambda \rightarrow \Delta$ be coverings with $\text{rk } \Lambda < \infty$, $\Psi \rightarrow \Delta$ Galois, and Ψ not simply connected. Then $\deg(\Lambda \rightarrow \Delta) < \infty$.*

The idea of the proof is that if the degree is infinite, then Λ has a hanging tree in its spine decomposition, and so Ψ does too. But Ψ should look the same at every point, hence *is* a tree.

Proof. Apply lattice excision to $\mathcal{L}(\Psi, \Delta)$, and as $\pi_1(\Psi, u)$ is unaffected by the excision of trees, we may assume that Δ is single vertexed. As $\deg(\Lambda \rightarrow \Delta)$ is infinite, the spine decomposition for Λ has an infinite tree, and Λ has the form of Proposition 8. Thus Ψ does too, by part (iii) of this Proposition, with subgraphs $\mathcal{Y}'_i \hookrightarrow \Psi$, edge e' and Ψ'_1 a tree. Take a closed reduced path γ in \mathcal{Y}'_2 , and choose a vertex u_1 of \mathcal{Y}'_1 such that the reduced path from u_1 to $s(e')$ has at least as many edges as γ . Project γ via the covering $\Psi \rightarrow \Delta$ to a closed reduced path, and then lift to u_1 . The result is reduced, closed by Proposition 14, and entirely contained in the tree \mathcal{Y}'_1 , hence trivial. Thus γ is also trivial so that \mathcal{Y}'_2 is a tree and Ψ is simply connected. \square

Proposition 21. *Let $\Lambda_u \rightarrow \Delta_v$ be a covering with $\text{rk } \Lambda < \infty$ and γ a non-trivial reduced closed path at v lifting to a non-closed path at u . Then there is an intermediate covering $\Lambda_u \rightarrow \Gamma_w \rightarrow \Delta_v$ with $\deg(\Gamma \rightarrow \Delta)$ finite and γ lifting to a non-closed path at w .*

Stallings shows something very similar [20, Theorem 6.1] starting from a finite immersion rather than a covering. As the proof shows, the path γ in Proposition 21 can be replaced by finitely many such paths. Moreover, the intermediate Γ constructed has the property that any set of Schreier generators for $\pi_1(\Lambda, u)$ can be extended to a set of Schreier generators for $\pi_1(\Gamma, w)$.

Proof. If $T \hookrightarrow \Delta$ is a spanning tree and $q : \Delta \rightarrow \Delta/T$ then γ cannot be contained in T , and so $q(\gamma)$ is non-trivial, closed and reduced. If the lift of $q(\gamma)$ to Λ/T_i is closed then the lift of γ to Λ has start and finish vertices that lie in the same component T_i of $p^{-1}(T)$, mapped homeomorphically onto T by the covering, and thus implying that γ is not closed. Thus we may apply lattice excision and pass to the single vertexed case while maintaining γ and its properties. Moreover, the conclusion in this case gives the result in general as closed paths go to closed paths when excising trees.

If the lift γ_1 of γ at u is not contained in the spine $\hat{\Lambda}_u$, then its terminal vertex lies in a tree T_{e_i} of the spine decomposition (\ddagger). By adding an edge if necessary to $\hat{\Lambda}_u \cup \gamma_1$, we obtain a finite subgraph whose coboundary edges are paired, with the edges in each pair covering the same edge in Δ , as below left:

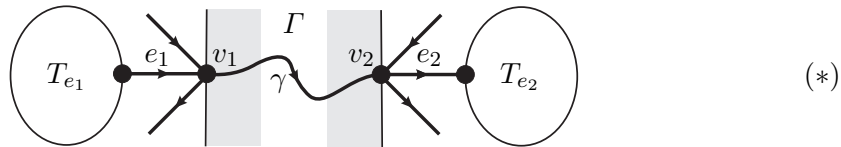


(if the lift is contained in the spine, take $\hat{\Lambda}_u$ itself). In any case, let Γ be $\hat{\Lambda}_u \cup \gamma_1$ together with a single edge replacing each pair as above right. Restricting the covering $\Lambda \rightarrow \Delta$ to $\hat{\Lambda}_u \cup \gamma_1$ and mapping the new edges to the common image of the old edge pairs gives a finite covering $\Gamma \rightarrow \Delta$ and hence by Lemma 3(i) an intermediate covering $\Lambda \xrightarrow{q} \Gamma \rightarrow \Delta$, with $q(\gamma_1)$ non-closed at $q(u)$. \square

For the rest of this section we investigate the rank implications of the decomposition (1) and the pairing (\ddagger) in a special case. Suppose $\Lambda \rightarrow \Delta$ is a covering with Δ single vertexed, $\text{rk } \Delta = 2$, Λ non-simply connected and $\text{rk } \Lambda < \infty$. Let $x_i^{\pm 1}$, ($1 \leq i \leq 2$) be the edge loops of Δ and fix a spine so we have the decomposition (1).

An *extended spine* for such a Λ is a connected subgraph $\Gamma \hookrightarrow \Lambda$ obtained by adding finitely many edges to a spine, so that every vertex of Γ is incident with either zero or three edges in its coboundary $\delta\Gamma$. It is always possible to find an extended spine: take the union of the spine $\hat{\Lambda}_u$ and each edge $e \in \delta\hat{\Lambda}_u$ in its coboundary. Observe that Γ is finite and the decomposition (1) gives $\text{rk } \Gamma = \text{rk } \hat{\Lambda}_u = \text{rk } \Lambda$. Call a vertex of the extended spine Γ *interior* (respectively *boundary*) when it is incident with zero (resp. three) edges in $\delta\Gamma$.

We have the pairing of trees (\ddagger) for an extended spine, so that each boundary vertex v_1 is paired with another v_2 ,



with e_1, e_2 and all the edges in the path $\gamma = \alpha(\mathcal{R}) \cap \Gamma$ covering an edge loop $x_i \in \Delta$. Call this an x_i -pair, ($i = 1, 2$).

For two x_i -pairs (fixed i), the respective γ paths share no vertices in common, for otherwise there would be two distinct edges covering the same $x_i \in \Delta$ starting at such a common vertex. Moreover, γ must contain vertices of Γ apart from the two boundary vertices v_1, v_2 , otherwise Λ would be simply connected. These other vertices are incident with at least two edges of $\gamma \in \Gamma$, hence at most 2 edges of the coboundary $\delta\Gamma$, and thus must be interior.

Lemma 8. *If $n_i, (i = 1, 2)$, is the number of x_i -pairs in an extended spine Γ , then the number of interior vertices is at least $\sum n_i$.*

(Lemma 8 is not true in the case $\text{rk } \Delta > 2$).

Proof. The number of interior vertices is $|V_\Gamma| - 2 \sum n_i$ and the number of edges of Γ is $4(|V_\Gamma| - 2 \sum n_i) + 2 \sum n_i$, hence $\text{rk } \Gamma - 1 = |V_\Gamma| - 3 \sum n_i$ by §1.6. As Λ is not simply connected, $\text{rk } \Lambda - 1 = \text{rk } \Gamma - 1 \geq 0$, thus $|V_\Gamma| - 2 \sum n_i \geq \sum n_i$ as required. \square

It will be helpful in §4 to have a pictorial description of the quantity $\text{rk} - 1$ for our graphs. To this end, a *checker* is a small plastic disk, as used in the eponymous boardgame (called *draughts* in British English). We place black checkers on some of the vertices of an extended spine Γ according to the following scheme: place black checkers on all the interior vertices of Γ ; for each x_1 -pair in $(*)$, take the interior vertex on the path γ that is closest to v_1 (ie: is the terminal vertex of the edge of γ whose start vertex is v_1) and *remove* its checker; for each x_2 -pair, we can find, by Lemma 8, an interior vertex with a checker still on it. Choose such a vertex and remove its checker also.

Lemma 9. *With black checkers placed on the vertices of an extended spine for Λ as above, the number of black checkers is $\text{rk } \Lambda - 1$.*

Proof. We saw in the proof of Lemma 8 that $\text{rk } \Lambda - 1 = \text{rk } \Gamma - 1$ is equal to the number of interior vertices of Γ less the number of x_i -pairs ($i = 1, 2$). \square

From now on we will only use the extended spine obtained by adding the coboundary edges to some fixed spine $\hat{\Lambda}_u$.

Let $p : \Lambda_u \rightarrow \Delta_v$ be a covering with $\text{rk } \Delta = 2$, $\text{rk } \Lambda < \infty$ and Λ not simply connected. A spanning tree $T \hookrightarrow \Delta$ induces a covering $\Lambda/T_i \rightarrow \Delta/T$ with Δ/T single vertexed. Let $\mathcal{H}(\Lambda_u \rightarrow \Delta_v)$ be the number of vertices of the spine of Λ/T_i at $q(u)$ and $n_i(\Lambda_u \rightarrow \Delta_v)$ the number of x_i -pairs in the extended spine. The homeomorphism class of Λ/T_i and the spine are independent of the spanning tree T , hence the quantities $\mathcal{H}(\Lambda_u \rightarrow \Delta_v)$ and $n_i(\Lambda_u \rightarrow \Delta_v)$ are too.

4. Pullbacks

Let $p_i : \Lambda_i := \Lambda_{u_i} \rightarrow \Delta_v, (i = 1, 2)$ be coverings and $(\Lambda_1 \prod_\Delta \Lambda_2)$ their (unpointed) pullback. If $\hat{\Lambda}_{u_i}$ is the spine at u_i then we can restrict the coverings to maps $p_i : \hat{\Lambda}_{u_i} \rightarrow \Delta_v$ and form the pullback $\hat{\Lambda}_{u_1} \prod_\Delta \hat{\Lambda}_{u_2}$.

Proposition 22 (spine decomposition of pullbacks). *The pullback $\Lambda = (\Lambda_1 \prod_\Delta \Lambda_2)$ has a wedge sum decomposition $\Lambda = (\hat{\Lambda}_{u_1} \prod_\Delta \hat{\Lambda}_{u_2}) \vee_\Theta \Phi$ with Φ a forest and no two vertices of the image of $\Theta \hookrightarrow \Phi$ lying in the same component.*

Proof. Let $\Lambda_i = \hat{\Lambda}_{u_i} \vee_{\Theta_i} \Phi_i, (i = 1, 2)$ be the spine decomposition, $t_i : \Lambda_1 \prod_\Delta \Lambda_2 \rightarrow \Lambda_i, (i = 1, 2)$ the coverings provided by the pullback and Ω a connected component of the pullback. If $\Omega \cap (\hat{\Lambda}_{u_1} \prod_\Delta \hat{\Lambda}_{u_2}) = \emptyset$, then a reduced closed path $\gamma \in \Omega$ must map via one of the t_i to a closed path in the forest Φ_i . As the images under coverings of reduced paths are reduced, $t_i(\gamma)$ must contain a spur which can be lifted to a spur in γ . Thus Ω is a tree.

Otherwise choose a vertex $w_1 \times w_2$ in $\Omega \cap (\hat{\Lambda}_{u_1} \prod_\Delta \hat{\Lambda}_{u_2})$ and let Γ be the connected component of this intersection containing $w_1 \times w_2$. If γ a reduced closed path at $w_1 \times w_2$ then $t_i(\gamma), (i = 1, 2)$ a reduced closed path at $w_i \in \hat{\Lambda}_{u_i}$, hence by Lemma 1(ii), $t_i(\gamma) \in \hat{\Lambda}_{u_i}$ and thus $\gamma \in \hat{\Lambda}_{u_1} \prod_\Delta \hat{\Lambda}_{u_2}$. Applying Proposition 18, we have Ω a wedge sum of Γ and a forest of the required form. \square

Corollary 5 (Howsen-Stallings). *Let $p_i : \Lambda_i \rightarrow \Delta, (i = 1, 2)$, be coverings with $\text{rk } \Lambda_i < \infty$ and $u_1 \times u_2$ a vertex of their pullback. Then $\text{rk } (\Lambda_1 \prod_\Delta \Lambda_2)_{u_1 \times u_2} < \infty$.*

Proof. The component Ω of the pullback containing $u_1 \times u_2$ is either a tree or the wedge sum of a finite graph and a forest as described in Proposition 22. Either case gives the result. \square

The remainder of this section is devoted to a proof of an estimate for the rank of the pullback of finite rank graphs in a special case. Let $p_j : \Lambda_j := \Lambda_{u_j} \rightarrow \Delta_v, (j = 1, 2)$ be coverings with $\text{rk } \Delta = 2, \text{rk } \Lambda_j < \infty$ and the Λ_j not simply connected. Let $\mathcal{H}_j := \mathcal{H}(\Lambda_{u_j} \rightarrow \Delta_v)$ and $n_{ji} := n_i(\Lambda_{u_j} \rightarrow \Delta_v)$ be as at the end of §3.

Theorem 4. For $i = 1, 2$,

$$\sum_{\Omega} (\text{rk } \Omega - 1) \leq \prod_j (\text{rk } \Lambda_j - 1) + \mathcal{H}_1 \mathcal{H}_2 - (\mathcal{H}_1 - n_{1i})(\mathcal{H}_2 - n_{2i}),$$

the sum over all non simply connected components Ω of the pullback $\Lambda_1 \prod_{\Delta} \Lambda_2$.

Proof. Lattice excision and the definition of the \mathcal{H}_j and n_{ji} allow us to pass to the Δ single vertexed case. Suppose then that Δ has edge loops $x_i^{\pm 1}, (1 \leq i \leq 2)$ at the vertex v , extended spines $\hat{\Lambda}_{u_j} \hookrightarrow \Gamma_j \hookrightarrow \Lambda_j$, and by restricting the covering maps p_j appropriately, the pullbacks $\hat{\Lambda}_{u_1} \prod_{\Delta} \hat{\Lambda}_{u_2} \hookrightarrow \Gamma_1 \prod_{\Delta} \Gamma_2 \hookrightarrow \Lambda_1 \prod_{\Delta} \Lambda_2$ with $t_j : \Lambda_1 \prod_{\Delta} \Lambda_2 \rightarrow \Lambda_j$ the resulting covering maps.

Place black checkers on the vertices of the extended spines Γ_j as in §3 and place a black checker on a vertex $v_1 \times v_2$ of $\Gamma_1 \prod_{\Delta} \Gamma_2$ precisely when both $t_j(v_j) \in \Gamma_j, (j = 1, 2)$ have black checkers on them. By Lemma 9, and the construction of the pullback for Δ single vertexed, we get the number of vertices in $\Gamma_1 \prod_{\Delta} \Gamma_2$ with black checkers is equal to $\prod (\text{rk } \Lambda_j - 1)$.

Let Ω be a non simply connected component of the pullback $\Lambda_1 \prod_{\Delta} \Lambda_2$ and $\mathcal{Y} = \Omega \cap (\Gamma_1 \prod_{\Delta} \Gamma_2)$. If $v_1 \times v_2$ is the start vertex of at least one edge in the coboundary $\delta \mathcal{Y}$, then at least one of the v_j must be incident with at least one, hence three, edges of the coboundary $\delta \Gamma_j$. Lifting these three via the covering t_j to $v_1 \times v_2$ gives at least three edges starting at $v_1 \times v_2$ in the coboundary $\delta \mathcal{Y}$. Four coboundary edges starting here would mean that Ω was simply connected, hence every vertex of \mathcal{Y} is incident with either zero or three coboundary edges.

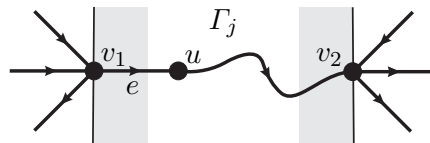
We can thus extend the interior/boundary terminology of §3 to the vertices of \mathcal{Y} , and observe that a vertex of \mathcal{Y} covering, via either of the t_j , a boundary vertex $\in \Gamma_j$, must itself be a boundary vertex. The upshot is that \mathcal{Y} is an extended spine in Ω and by Proposition 22, $\text{rk } \Omega - 1 = \text{rk } \mathcal{Y} - 1$. Now place *red* checkers on the vertices of \mathcal{Y} as in §3 and do this for each non-simply connected component Ω . The number of red checkered vertices is $\sum_{\Omega} (\text{rk } \Omega - 1)$.

The result is that $\Gamma_1 \prod_{\Delta} \Gamma_2$ has vertices with black checkers, vertices with red checkers, vertices with red checkers sitting on top of black checkers, and vertices that are completely unchecked. Thus,

$$\sum_{\Omega} (\text{rk } \Omega - 1) \leq \prod (\text{rk } \Lambda_j - 1) + N,$$

where N is the number of vertices of $\Gamma_1 \prod_{\Delta} \Gamma_2$ that have a red checker but no black checker.

It remains then to estimate the number of these “isolated” red checkers. Observe that a vertex of $\Gamma_1 \prod_{\Delta} \Gamma_2$ has no black checker precisely when it lies in the fiber, via at least one of the τ_j , of a checkerless vertex in Γ_j . Turning it around, we investigate the fibers of the checkerless vertices of both Γ_j . Indeed, in an x_1 -pair,



the vertices v_1, v_2 and u are checkerless, while v_1, v_2 are also checkerless in an x_2 -pair. We claim that no vertex in the fiber, via t_j , of these five has a red checker. A vertex of \mathcal{Y} in the fiber of the boundary vertices v_1, v_2 is itself a boundary vertex, hence contains no red checker. If $v_1 \times v_2 \in \mathcal{Y}$ is in the fiber of u and is a boundary vertex of \mathcal{Y} then it carries no red checker either.

If instead $v_1 \times v_2$ is an interior vertex then the lift to $v_1 \times v_2$ of e^{-1} cannot be in the coboundary $\delta\mathcal{Y}$, hence the terminal vertex of this lift is $\in \mathcal{Y}$ also and covers v_1 . Thus, this terminal vertex is a boundary vertex for an x_1 -pair of \mathcal{Y} , and $v_1 \times v_2$ is the interior vertex from which a red checker is removed for this pair.

The only remaining checkerless vertices of the Γ_j unaccounted for are those interior vertices chosen for each x_2 -pair, and thus $N \leq$ the number of vertices of $\Gamma_1 \prod_{\Delta} \Gamma_2$ contained in the fibers of these. If $u \in \Gamma_1$ is one of these interior vertices, then $u \times V_{\Gamma_2}$ are the vertices of $\Gamma_1 \prod_{\Delta} \Gamma_2$ in the fiber. As the boundary vertices in this fiber do not have red checkers we need only consider the $u \times \{\text{interior vertices of } \Gamma_2\}$ with these interior vertices precisely those of the spine $\widehat{\Lambda}_{u_2}$. Thus our fiber is $u \times \{\text{vertices of } \widehat{\Lambda}_{u_2}\}$, of which there are \mathcal{H}_2 , and a total of $n_{12}\mathcal{H}_2$ vertices of $\Gamma_1 \prod_{\Delta} \Gamma_2$ arising this way. There are also $n_{22}\mathcal{H}_1$ vertices arising in this way from $u \in \Gamma_2$, and $n_{12}n_{22}$ vertices counted twice.

Thus $N \leq n_{12}\mathcal{H}_2 + n_{22}\mathcal{H}_1 - n_{12}n_{22}$, hence the result for $i = 2$. Interchanging the checking scheme for the x_i -pairs gives the result for $i = 1$. \square

5. Free groups and the topological dictionary

A group F is *free of rank* $\text{rk } F$ if and only if it is isomorphic to the fundamental group of a connected graph of rank $\text{rk } F$. If Γ_1, Γ_2 are connected graphs with $\pi_1(\Gamma_1, v_1) \cong \pi_1(\Gamma_2, v_2)$, then we have $H_1(\Gamma_1) \cong H_1(\Gamma_2)$ and thus $\text{rk } \Gamma_1 = \text{rk } \Gamma_2$.

The free groups so defined are of course the standard free groups and the rank is the usual rank of a free group. At this stage we appeal to the existing (algebraic) theory of free groups, and in particular, that by applying Nielsen transformations, a set of generators for a free group can be transformed into a set of free generators whose cardinality is no greater. Thus, a finitely generated free group has finite rank (the converse being obvious). From now on we use the (topologically more tractable) notion of finite rank as a synonym for finitely generated.

Let F be a free group with representation $\varphi : F \rightarrow \pi_1(\Delta, v)$ for Δ connected. The topological dictionary is the loose term used to describe the correspondence between algebraic properties of F and topological properties of Δ as described in §§1-2. The non-abelian F correspond to the Δ with $\text{rk } \Delta > 1$. A subgroup $A \subset F$ corresponds to a covering $p : \Lambda_u \rightarrow \Delta_v$ with $p^*\pi_1(\Lambda, u) = \varphi(A)$, and hence $\text{rk } A = \text{rk } \Lambda$. Thus finitely generated subgroups correspond to finite rank Λ and normal subgroups to Galois coverings. Inclusion relations between subgroups correspond to covering relations, indices of subgroups to degrees of coverings, trivial subgroups to simply connected coverings, conjugation to change of basepoint.

Applying the topological dictionary to the italicised results below we recover some classical facts (see also [19, 20]).

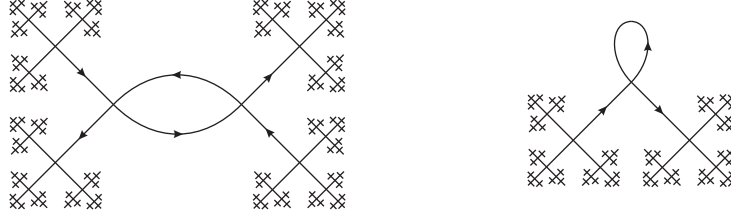
1. [7, 12]: If a finitely generated subgroup A of a non-abelian free group F is contained in no subgroup of infinite rank, then A has finite index in F ; *Proposition 19*.
2. [7]: If a finitely generated subgroup A of a free group F contains a non-trivial normal subgroup of F , then it has finite index in F ; *Proposition 20*.
3. [1, 8]: Let F be a free group, X a finite subset of F , and A a finitely generated subgroup of F disjoint from X . Then A is a free factor of a group G , of finite index in F and disjoint from X ; *Proposition 21* (and the comments following it).
4. [9]: If A_1, A_2 are finitely generated subgroups of a free group F , then the intersection of conjugates $A_1^{g_1} \cap A_2^{g_2}$ is finitely generated for any $g_1, g_2 \in F$; *Corollary 5*.

If Δ is a graph, $\text{rk } \Delta = 2$, and $A \subset F = \pi_1(\Delta, v)$, then we define $\mathcal{H}(F, A) := \mathcal{H}(\Lambda_u \rightarrow \Delta_v)$ and $n_i(F, A) := n_i(\Lambda_u \rightarrow \Delta_v)$, where $p : \Lambda_u \rightarrow \Delta_v$ is the covering with $p^*\pi_1(\Lambda, u) = A$. For an arbitrary free group F with representation $\varphi : F \rightarrow \pi_1(\Delta, v)$, define $\mathcal{H}^\varphi(F, A)$ and $n_i^\varphi(F, A)$ to be $\mathcal{H}(\varphi(F), \varphi(A))$ and $n_i(\varphi(F), \varphi(A))$.

The appearance of φ in the notation is meant to indicate that these quantities, unlike rank, are representation dependent. This can be both a strength and a weakness. A weakness because it seems desirable for algebraic statements to involve only algebraic invariants, and a strength if

we have the freedom to choose the representation, especially if the most interesting results are obtained when this representation is not the “obvious” one.

For example, if F is a free group with free generators x and y , and Δ is single vertexed with two edge loops whose homotopy classes are a and b , then the subgroup $A = \langle xy \rangle \subset F$ corresponds to the Λ below left under the obvious representation $\varphi_1(x) = a, \varphi_1(y) = b$, and to the righthand graph via $\varphi_2(x) = a, \varphi_2(y) = a^{-1}b$:



Thus, $\mathcal{H}^{\varphi_1}(F, A) = 2, n_i^{\varphi_1}(F, A) = 1, (i = 1, 2)$, whereas $\mathcal{H}^{\varphi_2}(F, A) = 1, n_1^{\varphi_2}(F, A) = 1, n_2^{\varphi_2}(F, A) = 0$.

We now apply the topological dictionary to Theorem 4. Let $\varphi : F \rightarrow \pi_1(\Delta, v)$, $A_j \subset F$, ($j = 1, 2$), finitely generated non-trivial subgroups, and $p_j : \Lambda_{u_j} \rightarrow \Delta_v$, ($j = 1, 2$) coverings with $\varphi(A_j) = p_j^* \pi_1(\Lambda, u_j)$. Each non simply-connected component Ω of the pullback corresponds to some non-trivial intersection of conjugates $A_1^{g_1} \cap A_2^{g_2}$. As observed in [13], these in turn correspond to the conjugates $A_1 \cap A_2^g$ for g from a set of double coset representatives for $A_2 \backslash F / A_1$.

Theorem 5. *Let F be a free group of rank two and $A_j \subset F$, ($j = 1, 2$), finitely generated non-trivial subgroups. Then for any representation $\varphi : F \rightarrow \pi_1(\Delta, v)$ and $i = 1, 2$,*

$$\sum_g (\text{rk}(A_1 \cap A_2^g) - 1) \leq \prod_j (\text{rk } A_j - 1) + \mathcal{H}_1 \mathcal{H}_2 - (\mathcal{H}_1 - n_{1i})(\mathcal{H}_2 - n_{2i}),$$

the sum over all double coset representatives g for $A_2 \backslash F / A_1$ with $A_1 \cap A_2^g$ non-trivial, and where $\mathcal{H}_j = \mathcal{H}^\varphi(F, A_j)$ and $n_{ji} = n_i^\varphi(F, A_j)$.

This theorem should be viewed in the context of attempts to prove the so-called *strengthened Hanna Neumann conjecture*: namely, if A_j , ($j = 1, 2$) are finitely generated, non-trivial, subgroups of an arbitrary free group F , then

$$\sum_g (\text{rk}(A_1 \cap A_2^g) - 1) \leq \prod_j (\text{rk } A_j - 1) + \varepsilon,$$

the sum over all double coset representatives g for $A_2 \backslash F / A_1$ with $A_1 \cap A_2^g$ non-trivial, where the conjecture is that ε is zero, while in the existing results, it is an error term having a long history. We provide a very partial, and chronological, summary of these estimates for ε in the table:

$(\text{rk } A_1 - 1)(\text{rk } A_2 - 1)$	H. Neumann [14]
$\max\{(\text{rk } A_1 - 2)(\text{rk } A_2 - 1), (\text{rk } A_1 - 1)(\text{rk } A_2 - 2)\}$	Burns [1]
$\max\{(\text{rk } A_1 - 2)(\text{rk } A_2 - 2) - 1, 0\}$	Tardos [21]
$(\text{rk } A_1 - 3)(\text{rk } A_2 - 3)$	Dicks-Formanek [4]

(the original, unstrengthened, conjecture [14] involved just the intersection of the two subgroups, rather than their conjugates, and the first two expressions for ε were proved in this restricted sense; the strengthened version was formulated in [13], and the H. Neumann and Burns estimates for ε were improved to the strengthened case there). Observe that as the join $\langle A_1, A_2 \rangle$ of two finitely generated subgroups is finitely generated, and every finitely generated free group can be embedded as a subgroup of the free group of rank two, we may replace the ambient free group in the conjecture with the free group of rank two.

It is hard to make a precise comparison between the ε provided by Theorem 5 and those in the table. Observe that if $A_j \subset F$, with F free of rank two, then with respect to a topological representation we have $\text{rk } A_j = \mathcal{H}_j - (n_{j1} + n_{j2}) + 1$. It is straight forward to find infinite families $A_{1k}, A_{2k} \subset \pi_1(\Delta, v)$, ($k \in \mathbb{Z}^{>0}$), for which the error term in Theorem 5 is less than those in the table above for all but finitely many k , or even for which the strengthened Hanna Neumann conjecture is true by Theorem 5, for instance,

$$A_{1k} = A_{2k} = \begin{array}{c} \text{Diagram showing a graph with two vertices connected by a horizontal edge. Each vertex has four incident edges forming a cross shape. Between the two vertices, there are k loops (represented by circles) on the horizontal edge. The diagram is labeled 'k edge loops' below the loops. } \end{array}$$

but where the error terms in the table are quadratic in k .

References

- [1] R. G. Burns, *A note on free groups*, Proc. Amer. Math. Soc. **23** (1969), 14–17.
- [2] Daniel E. Cohen, *Combinatorial group theory: a topological approach*, London Mathematical Society Student Texts, vol. 14, Cambridge University Press, Cambridge, 1989, ISBN 0-521-34133-7, 0-521-34936-2.
- [3] D. J. Collins and H. Zieschang, *Combinatorial group theory and fundamental groups* (1993), 1–166, 233–240.
- [4] Warren Dicks and Edward Formanek, *The rank three case of the Hanna Neumann conjecture*, J. Group Theory **4** (2001), 113–151.
- [5] Warren Dicks, *Equivalence of the strengthened Hanna Neumann conjecture and the amalgamated graph conjecture*, Invent. Math. **117** (1994), 373–389.
- [6] S. M. Gersten, *Intersections of finitely generated subgroups of free groups and resolutions of graphs*, Invent. Math. **71** (1983), 567–591.
- [7] Leon Greenberg, *Discrete groups of motions*, Canad. J. Math. **12** (1960), 415–426.
- [8] Marshall Hall Jr., *Subgroups of finite index in free groups*, Canadian J. Math. **1** (1949), 187–190.
- [9] A. G. Howson, *On the intersection of finitely generated free groups*, J. London Math. Soc. **29** (1954), 428–434.
- [10] Wilfried Imrich, *On finitely generated subgroups of free groups*, Arch. Math. (Basel) **28** (1977), 21–24.
- [11] ———, *Subgroup theorems and graphs* (1977), 1–27. Lecture Notes in Math., Vol. 622.
- [12] Abraham Karrass and Donald Solitar, *On finitely generated subgroups of a free group*, Proc. Amer. Math. Soc. **22** (1969), 209–213.
- [13] Walter D. Neumann, *On intersections of finitely generated subgroups of free groups* (1990), 161–170.
- [14] Hanna Neumann, *On the intersection of finitely generated free groups*, Publ. Math. Debrecen **4** (1956), 186–189.
- [15] Peter Nickolas, *Intersections of finitely generated free groups*, Bull. Austral. Math. Soc. **31** (1985), 339–348.
- [16] K Reidemeister, *Fundamentalgruppen und Überlagerungsräume*, Nachr. Ges. Wiss. Göttingen, Math. Phys. Kl. (1928), 69–76. (German)
- [17] Peter Scott and Terry Wall, *Topological methods in group theory*, Homological Group Theory (Proc. Sympos., Durham, 1977), London Math. Soc. Lecture Note Ser., vol. 36, Cambridge Univ. Press, Cambridge, 1979, pp. 137–203.
- [18] Jean-Pierre Serre, *Trees*, Springer Monographs in Mathematics, Springer-Verlag, Berlin, 2003, ISBN 3-540-44237-5, Translated from the French original by John Stillwell; Corrected 2nd printing of the 1980 English translation.
- [19] Brigitte Servatius, *A short proof of a theorem of Burns*, Math. Z. **184** (1983), 133–137.
- [20] John R. Stallings, *Topology of finite graphs*, Invent. Math. **71** (1983), 551–565.
- [21] Gábor Tardos, *Towards the Hanna Neumann conjecture using Dicks’ method*, Invent. Math. **123** (1996), 95–104.
- [22] ———, *On the intersection of subgroups of a free group*, Invent. Math. **108** (1992), 29–36.